AD 601641

SEQUENTIAL SEARCH AND DETECTION

by

Stephen Michael Pollock

Technical Report No. 5
On Work Performed Under
Contract Nonr-3963 (06), Office of Naval Research

Applications of Probabilistic System Models
to Naval Problems

Operations Research Center
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
Cambridge 39, Massachusetts

May, 1964

*Adapted from a thesis presented to the Department of Physics in partial fulfillment of the requirements for the degree of Doctor of Philosophy, May 1964.

FOREWORD

The Center for Operations Research at the Massachusetts Institute of Technology is an inter-departmental activity devoted to graduate training and research in the field of operations research. Its products are books, journal articles, detailed reports such as this one, and students trained in the theory and practice of operations research.

The work of the Center is supported, in part, by government contracts and industrial grants-in-aid. The research which led to this particular report was supported by the Office of Naval Research under Contract Nonr-3963 (06).

Philip M. Morse,
Director of the Center

Ronald A. Howard,
Project Supervisor

ACKNOWLEDGEMENT

The author wishes to express his appreciation to the many people who have assisted in the conception and preparation of the work presented here.

Particular thanks are due to Professor Ronald A. Howard who supervised this study, and who was always available for interesting and constructive commentary. Professors Philip M. Morse and George P. Wadsworth rendered valuable aid by criticising and strengthening the final manuscript. Thanks are also due to Dr. George Murray, Mr. T. Roderburg and others at the M. I. T. Operations Research Center, for their helpful discussions and comments.

Finally, for the grace, speed and patience applied to the arduous task of typing the manuscript, acknowledgement is gratefully given to Mrs. Antoinette Jailiet.

This work was supported in part by the Office of Naval Research and the M. I. T. Department of Physics, who sponsored the author's Research Assistantship at the Operations Research Center. Part of this work was done at the M. I. T. Computation Center.

SEQUENTIAL SEARCH AND DETECTION

by

STEPHEN MICHAEL POLLOCK

Submitted to the Department of Physics on May 15, 1964 in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

ABSTRACT

The objective of this study has been to obtain and evaluate strategies to be used in certain general search situations. These strategies minimize the expected cost of search and resulting decisions and are sequential in the sense that a decision at any time is dependent upon what has been observed up to that time.

The first situation studied leads to the formulation of a minimum expected cost sequential hypothesis test. The target is either present in the region of interest with a priori probability P, or not with probability $1-P$. Knowing the value of P, at fixed intervals of time the searcher must either make a terminal decision (i.e. decide that the target is present, or not present) or make a measurement of a random variable that has a probability density function which depends upon whether or not the target is present. A cost structure is given which assigns costs to wrong terminal decisions, as well as a cost (which depends upon whether or not the target is present) for the taking of a measurement. The sequential strategy and resultant minimum cost are derived by solving a functional equation of the dynamic programming type. The relation between this strategy and the Wald sequential probability ratio test is discussed. The minimum cost of the strategy is compared with the cost of an often used non-sequential strategy as well as a class of sub-optimal sequential strategies that involve threshold observations.

The second part of this study involves a situation in which it is assumed that the target arrives at some random time (the "raid-recognition" problem). A cost structure is given which assigns a cost to deciding the target has arrived when in fact it hasn't, and also assigns a cost proportional to the time between arrival of the target and the decision that it has arrived. Again observations of a random variable related to the presence of the target are available as an alternative to making such a decision. The sequential strategy and resulting minimum cost are again obtained by means of a functional equation. An additional result is the formulation of a System Operating Characteristic that is used for this randomly arriving target model in a way similar to the use of the Receiver Operating Characteristic for the hypothesis test model.

Thesis Supervisor: Ronald A. Howard
Title: Associate Professor of Electrical Engineering and
Associate Professor of Industrial Management

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

CHAPTER I

INTRODUCTION

Much work has been done in the past 20 years on analytical approaches to problems dealing with search and detection. The purpose of the present work is to extend some of these problems and to treat them from the point of view of Statistical Decision Theory.

In particular, the often separately stated problems in "search" and "detection" will be considered as part of an overall operation, which must consequently be optimized as a whole.

Most people have found themselves in the position of having to look for something. The process of looking for and (possibly) finding the object is, in the sense of what will be treated in this paper, that person's solution to a search and detection problem. Whether that solution is a "good" one or not depends upon the criterion of goodness that the person has decided upon. Hopefully, if the person has a degree of consistency which we would like to require of decision makers, he could describe his scheme of operation to a confederate, and be confident that the subsequent results would be the same as if he himself carried out the process. In this age of automated decision making, the confederate is very often a computer, and the description of the scheme the appropriate program.

One object of this paper is to consider a class of decision problems that might be called "sequential search and detection" problems, and using a particular utility structure associated with these problems arrive at a description of the "optimal" way of making these decisions:

optimal in the sense of minimizing the overall cost of the search process. In addition we are interested in examining some non-optimal decision schemes, some existing and some proposed, that do not minimize cost, but which are perhaps easily implemented, or have intuitive appeal, or both.

Before proceeding further, it is convenient to define the terms "search" and "detection". As nouns (and adjectives) they have often appeared in the literature as synonyms, but for the purposes of this paper there will be a clear distinction made between them.

Detection involves the gathering of information pertaining to the object being sought (the target), the sifting out what is useful information and the relaying of this in some efficient form to a decision maker.

Search describes the decisions made on the basis of the detection information received. In particular, a "search strategy" will be that set of rules that associates decisions with every conceivable result of the detection process.

When a collar button falls to the floor, the detection device of the eye picks up the information that there is a strong glint of light under the bureau, and a less strong one under the bed. The search phase is the decision to bend down and reach under the bureau, and if unsuccessful, to then reach under the bed. To continue the homely example, the search also included the decision to bend down in the first place, rather than shrugging the shoulders and taking a new button.

The problems of search are decision problems. Where should one look, for how long, with what equipment? The answers to these questions involve decisions, and in that any old answer is not acceptable we recognize the fact that there are costs involved in doing the wrong things: making the wrong decisions. In addition we often must use detection devices that give information imperfectly, perhaps describable statistically. Thus statistical decision theory is the most likely candidate to be the tool with which to attack search and detection problems.

1.1 Some definitions

The definitions offered below are solely for use in the context of this work, and the author makes no claims for their universal appeal or application.

Target

The target is an object that is of primary interest to the decision maker (sometimes called the "searcher"). At the starting point of the search process, the location of the target is uncertain. The general object of the search process is to increase the searcher's knowledge of the location of the target. In general, the target may or may not change its location during the search. It will be assumed, for the purposes of this work, that only one target is involved at any time. This assumption, although often unrealistic, helps to point to good procedures even in the cases where it does not hold.

Field (F)

The field is the area within which the entire search process takes place: the region of interest of the search.

Cell (A_i)

In most searches, the field is broken down into many smaller non-overlapping areas, A_i, called cells. In general a cell is, in size and shape, the resolution element of the sensing device being used--- the smallest space that could contain two targets without allowing the detection device to determine whether or not there are one or more targets present. In the discussion that follows, $i=1, 2, \ldots, M$, where M is the number of cells in the field. The cell A_0 represents the location "nowhere in the field".

State of Nature (S_i)

The state of nature is a description of the actual location of the target. The abbreviation S_i stands for the state of nature: {the target is present in cell A_i}. The abbreviation S_0 stands for the state of nature: {the target is not present anywhere in F}.

A Priori Target Location Probability Vector (P)

This vector summarizes the degree of uncertainty concerning the target location at the start of the search, where P_i is the a priori probability that S_i is the state of nature. The author does not wish to embroil himself in arguments concerning the existence of this $\underline{P}$ vector To the unconvinced reader two suggestions are offered: he can consider the placement of the target to be determined by some appropriate randor experiment (e.g. dice throwing), the outcomes of which can be associate with the P_i's ; or he can replace the word "probability" by "plausabilit throughout the remainder of this work, with the assurance (for example see Jeffreys (23)) that the mathematical development will be identical).

Detection Process

The detection process is that process by which information (concerning the presence or absence of the target in one, some or all of the cells) is obtained by means of observation or measurements of various physical phenomena related to these cells; such as reflected energy, radiated energy, etc. The apparatus used for these measurements is called the "detection device", or the "receiver".

Noise

Noise is the collection of those factors (random or otherwise) that make the detection device produce target-like signals when the target is in fact absent.

Decision (D_i or W)

A decision (D_i) is a commitment by the searcher to take action associated with the belief that the target is in cell A_i. Or it may be the commitment to wait (W) for more information from the detection device. The decisions D_i are called "terminal decisions".

Search Strategy

A search strategy is a set of rules which assigns decisions to all possible outcomes of the detection process.

Right Decision

A right decision occurs when D_j is made and S_j is the state of nature, written $\{D_j | S_j\}$.

Wrong Decision

A wrong decision occurs when D_j is made and S_i $(i \neq j)$ is the state of nature, written $\{D_j | S_i\}$, $(i \neq j)$.

Detection

A detection is one of the right decisions $\{D_j | S_j\}$ with $j \neq 0$.

False Alarm

A false alarm is one of the wrong decisions $\{D_j | S_i\}$ with $i=0$.

The Symbol ":"

Throughout this study a colon ":" in a mathematical express·· represents the phrase "make decision". Thus the statement

"if $x > x^*$: D_1 "

is read

"if x is greater than x^*, make decision D_1 ."

Events of Probability Zero

When dealing with a continuous random variable x with pro·bility density function (abbreviated p. d. f.) $f(x)$, it shall be assumed fo simplicity of notation that the p. d. f. is such that for any x^*

$$\text{prob.}\{x \leq x^*\} = \text{prob.}\{x < x^*\} = 1 - \text{prob.}\{x \geq x^*\}$$

i.e. $\text{prob.}\{x = x^*\} = 0$ for all x^* .

CHAPTER II

BACKGROUND

2.1 General

Published work on the subject of search and detection seems to fall into three separate groups, with, as a rule, little or no discussion of the relation among them. Koopman (26) has categorized these groups as dealing with 1) Kinematics, 2) Distribution of Search Effort, 3) Target Detection.

Problems concerning "kinematics" have to do primarily with the relative motions between searcher and target. Detection is assumed perfect, once an interception occurs, and the problems are directed towards establishing optimal pursuit and evasion strategies, prediction of target course, etc. The results, which are rather complete and cover a wide range of models, are to be found primarily in Koopman (26, 27), as well as Gluss (14, 15), Danskin (11), Isbell (22), Banta (1) and others. These kinematic considerations are outside the scope of the present work, inasmuch as the problems of imperfect detection and possible false alarms and decisions are not involved.

The results of previous work in both Distribution of Search Effort and Target Detection set the groundwork for the present paper, and the following two sections develop what has been done in these areas.

2.2 Distribution of Search Effort

Most of the work published under the general topic of "sear theory" has been involved with Koopmans' second category: distribut of search effort. These problems in general assume that the availabl detection devices are ones that will not yield target-like information the target is not present and will detect only some fraction of the tin the target is present. This fraction is called the detection probabilit

Thus the search strategy becomes straightforward. Whe target-like information is received from A_i, the target must be the so decide D_i.

The problem thus becomes not what to do when target-like information is received, but what to do until it is received.

The solutions are in the form of search strategies, involv the distribution of the available "magnitude" of search effort that sh be placed upon the cells. It is assumed that relations between the amount of search effort used in observing a cell, and quality of detec tion in that cell (as reflected in detection probability) are known quant tively. The distribution of effort is selected to optimize some entire measure of the search, such as maximizing the probability of detecting the target for a given search effort. Another criterion (whi proves to yield identical strategies for most models) would be to mini the average search effort needed to eventually detect the target.

The most general solutions to this problem have been obtai when the field represents a continuum (the cells becoming differential areas), by Koopman (32, 33) and de Geunin (22). Other cases, involv

so called "discrete" or finite sized cell cases, have been treated by others (5, 6, 10, 13, 30, 35, 36, 37) with solutions all approaching Koopman's in the limit as cell sizes decrease. However, all these approaches, as pointed out above, neglect the possibility of, and hence the cost of false alarms due to noise.

2.3 Target Detection - STSD and Hypothesis Testing

What is known as the Statistical Theory of Signal Detection (STSD) has been developed in the past decade, primarily by Middleton (33, 34). Helstrom (19) and others have extended this theory.

Hypothesis testing is used as a framework upon which to develop receivers that take into consideration the presence of noise in each cell. The theory is quite elegant, and much of it is concerned with problems of devising measures (or statistics) of received waveforms, allowing for the possibility of noise and target signals of all statistical varieties. However, the theory tends to not differentiate between search and detection. It also has tended to ignore the potential of the use of sequential rules involving costs of wrong decisions. In order to make some of these points clear, as well as to lay the ground for a more general extension of the theory, let us consider the basic ideas involved in the STSD.

Consider the field to consist of only one cell, A_1. Thus, there are two possible states of nature:

S_0 = target is not present.

S_1 = target is present.

It is assumed that the target is stationary. That is, only one of these states describes the target position for the entire duration of the search. The target cannot move, or evade the searcher. (This, of course, limits the generality of the model, and the eventual lifting of this condition is one of the aims of this study.)

Now, upon observing the cell with the detection device, the searcher can make one of two decisions, which by the assumptions of STSD are equivalent to taking one of two actions:

D_0 = take action appropriate to target being absent.

D_1 = take action appropriate to target being present.

Let us represent the observation by x (possibly a vector).

If the observer knows from experience that the sensing device operates in such a way that if a target is present (S_1) the reading would be x_1, but if the target is not present (S_0), the reading would be x_0, then the decision is sraight forward:

$$\text{when } x = x_0 : D_0$$
$$x = x_1 : D_1 .$$

(Note that a reading of $x \neq x_0$ or x_1 is impossible.)

A more realistic measurement, however, is subject to random fluctuations due to many incomprehensible or unpredictable effects including those in the processing equipment, as well as target characteristics, errors of observation, etc. The outcome of the observation, x, under the two possible hypotheses, is thus a random variable, and is at best describable by known probability density functions. If these probability density functions are completely described and involve no unknown parameters then the hypotheses are called "simple". This more realistic example, then, is the testing of two simple hypotheses, where the outcome x will have a known probability density function $p_0(x)$ under hypothesis S_0, and $p_1(x)$ under hypothesis S_1.

The observer must be able to decide upon either D_0 or D_1 for any possible observed value of x. The resulting decision rule divides the range of x into two exhaustive regions X_0 and X_1, with the result:

$$\begin{aligned} &\text{if } x \in X_0 : D_0 \\ &\text{if } x \in X_1 : D_1 . \end{aligned} \tag{2.1}$$

Because of the probabilistic nature of the observation, it is now possible for the observer to make a mistake: to say one hypothesis is correct when in fact the other is actually true. A way to describe the possibility of a wrong decision is to associate with any decision rule such as (2.1) above two error probabilities. These are defined as follows:

If the true state of nature is described by S_0, then the probability of making decision D_1 (a mistake) is defined as α, or "error of the first kind". Written in the usual probabilistic notation, $\alpha = \text{prob.}\ \{D_1 | S_0\}$. Correspondingly, the probability of making D_0 when the true state of nature is in fact described by S_1 is defined by β or "error of the second kind". That is, $\beta = \text{prob.}\ \{D_0 | S_1\}$. In the usual statistical sense, S_0 is said to be the "null hypothesis". Our choice, in this detection model, of S_0 being "target absent" is completely a matter of convenience.

In the detection model that we are describing, it is now a conceptual advantage to speak of a detection probability p_d and a false alarm probability p_f. These are a matter of definition, and it is convenient here to relate them to α and β. Thus,

$$p_d = \text{prob.}\ \{D_1 | S_1\} = 1 - \text{prob.}\ \{D_0 | S_1\} = 1 - \beta$$

$$p_f = \text{prob.}\ \{D_1 | S_0\} = \alpha$$

(note that these probabilities are strictly conditional upon the state of nature, and that conclusions reached by consideration of these numbers alone must in themselves be conditional. This will be discussed in the next section.)

In terms of the known p.d.fs $p_0(x)$ and $p_1(x)$ we see that α and β are given by

$$\alpha = p_f = \int_{X_1} p_0(x)\,dx$$

$$\beta = 1 - p_d = \int_{X_0} p_1(x)\,dx \;.$$

(2.[illegible])

As an example, let us imagine that in the cell of interest,

$$p_0(x) = f_N(x;\ 0, \sigma)$$

$$p_1(x) = f_N(x;\ \mu, \sigma)$$

where

$$f_N(x;\ \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right]$$

then a possible decision rule would be to select a constant x^* such that (see Figure 2.1)

$$\text{for } x \leq x^* : D_0$$

$$x > x^* : D_1 \;.$$

Here $X_0 = [-\infty, x^*]$ and $X_1 = [x^*, \infty]$. The error probabilities would be then given by

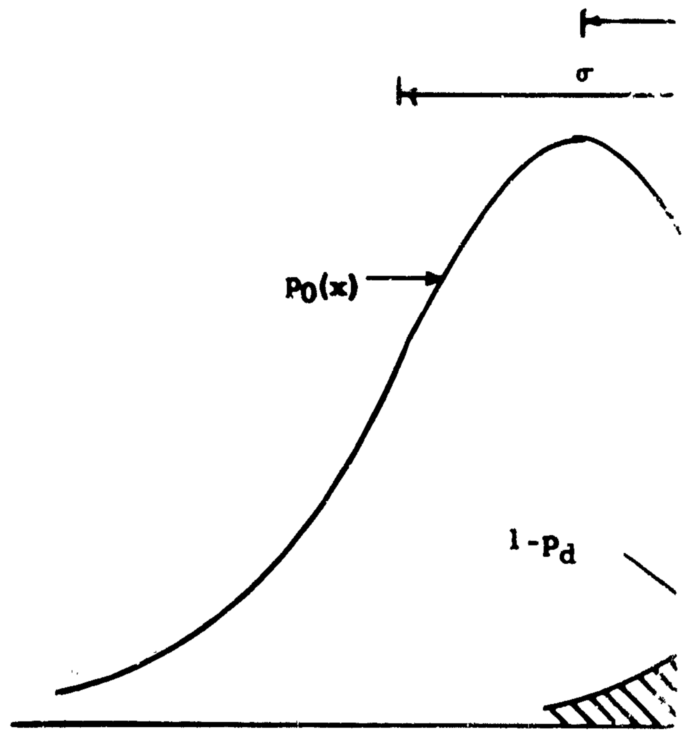

Figure 2.1 Example

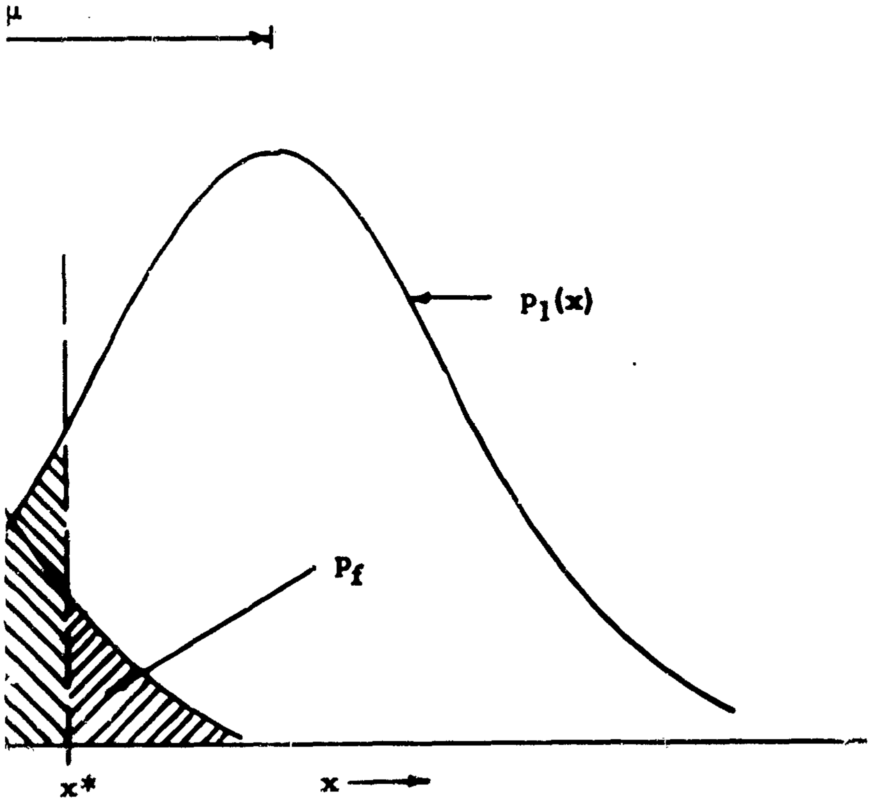

), p_1 (x) and threshold setting x*

$$\alpha = 1 - \mathrm{erf}\left(\frac{x^*}{\sigma}\right)$$

$$\beta = \mathrm{erf}\left(\frac{x^* - \mu}{\sigma}\right) .$$

For constant μ and σ, α and β are related to the parameter x^* so that as x^* increases, α decreases and β increases. This is shown in the curves in Figure 2.2, which are for $\sigma = 1$. Curves such as these, that relate α to β (or p_d to p_f) for a particular device are called ROC or Receiver Operating Characteristic curves. Figure 2.2 shows the ROC of p_d v.s. p_f with μ as a parameter. It is sometimes convenient to draw the ROC as a plot of p_d v.s. μ, with p_f as a parameter, where μ is in general some measure of the differ- between the p.d.f.s $p_0(x)$ and $p_1(x)$.

At this point we have not yet determined what the observer is actually trying to do. He might have some vague ideas about achiev- a "reasonable" detection probability while keeping the false alarm prob- bility "low". A quantification of this concept is the subject matter of the next section. However, at this point, we can agree upon some aspect of goodness, so that we can immediately compare certain sensing devices.

Suppose two devices are available to use in searching a cell. Device A has $p_d = 1-\beta = .8$, $p_f = \alpha = .01$, while device B has $p_d = 1-\beta = .8$, $p_f = \alpha = .02$. Clearly, device A is beter than device B by whatever reasonable criterion is used. This concept is formalize in the following way.

Let us consider two search devices d_1, d_2 with error probabilities (α_1, β_1) and (α_2, β_2) respectively. Then d_1 is

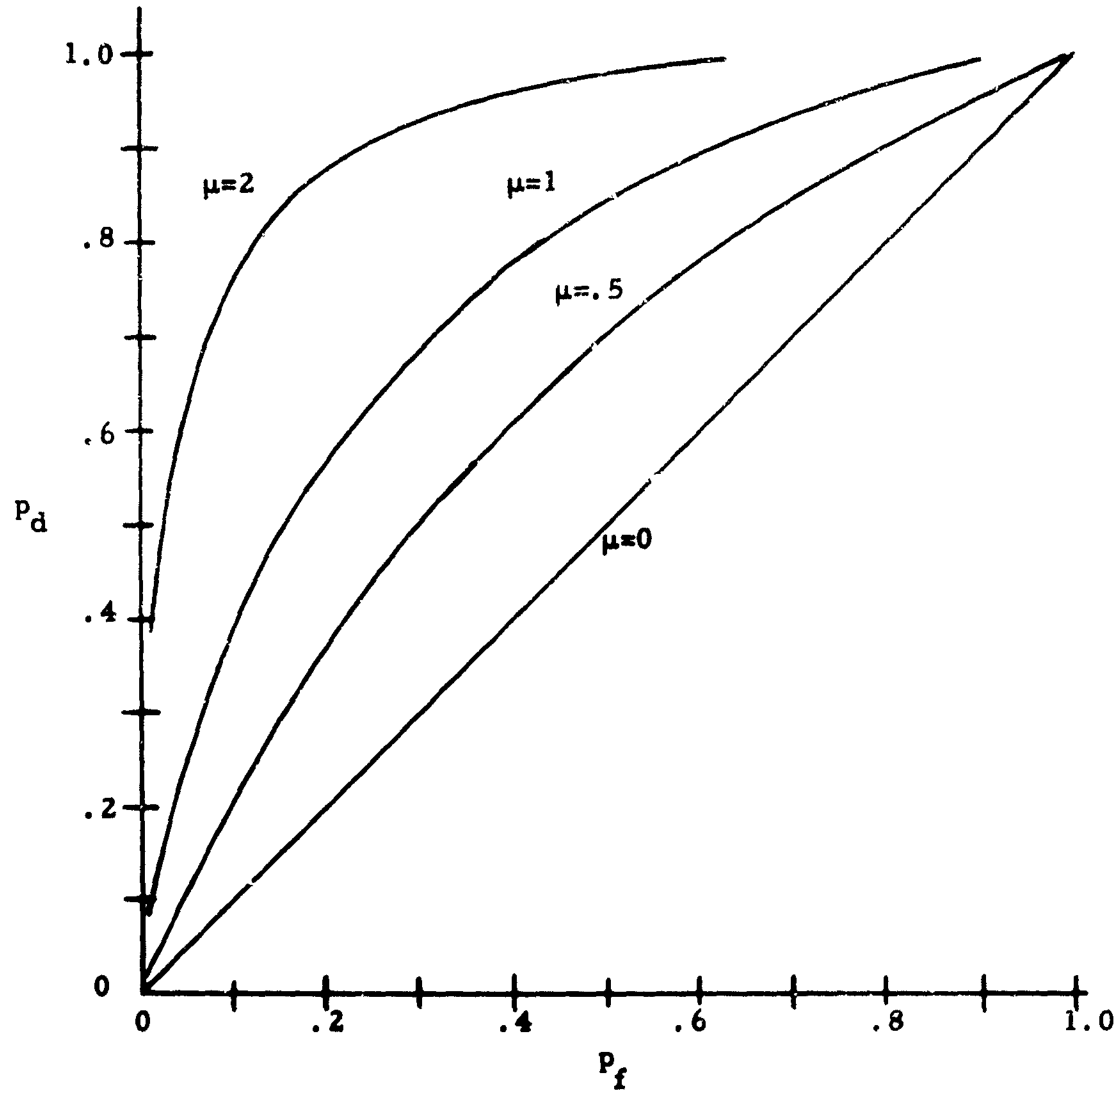

Figure 2.2 Receiver Operating Characteristic (ROC) for device illustrated in Figure 2.1

said to be preferable to, or better than, d_2 if $\alpha_1 \leq \alpha_2$ and $\beta_1 \leq \beta_2$. Thus, if all possible devices are represented on a p_d v.s. p_f ROC curve, any given device is better than those that are "below and to the right", while it is worse than those that are "above and to the left". However, if $\alpha_1 \leq \alpha_2$ and $\beta_1 \geq \beta_2$ then d_1 and d_2 are said to be non-comparable, and more must be known before a comparison may be made between the two.

Looking at the ROC in Figure 2.2, we see that for a given value of μ, as x^* is varied from minus infinity to infinity, a whole continuum of possible devices is described, all of these devices being non-comparable in the sense given above. For a given μ we cannot yet decide which among them is best. However, if we compare the curves for μ_1 and μ_2 ($\mu_2 > \mu_1$) we see that for any point (representing a device) on the μ_1 curve, we can find "above and to the left" of it some point on the μ_2 curve. Thus, an increase in μ is always desirable in the system as illustrated in Figure 2.2, as long as x^* is free to be set.

In the statistical literature, α is called the "size" of a particular test or experiment, and for a given value of α the largest attainable value of $1-\beta = p_d$ is called the "power".

2.4 Decision Criteria

In the last section, we have seen that a detection device combined with a decision rule may be operationally described by the associated error probabilities α and β. Let us call a combination of a detection device and a decision rule a "system".

In order to evaluate a system, or equivalently, to be able to compare two or more systems, there must exist some measure of the effectiveness of the entire system (including the specific decision rule used). The observer must be searching for the target with some idea as to not only what actions will be taken, but the costs associated with them, if he detects (or thinks he detects), or does not detect the target.

The object is now to determine which decision rule is "best" in some way. For example, in the system illustrated in Figure 2.1 what value of threshold x* should be used?

Since it is reasonable to assume that the search system will be used over and over again for some length of time, systems may be compared in terms of their overall, long-term average performance.

For each state of nature, the cost of making decision D_0 or D_1 will be different. In general, it is possible to list the costs of all possible state---decision situations in a matrix C:

$$C = \begin{bmatrix} C_{00} & C_{01} \\ C_{10} & C_{11} \end{bmatrix}$$

where C_{ij} is the cost of making decision D_i when in fact S_j is true. (For this to be a reasonable matrix $C_{00} < C_{01}$; $C_{11} < C_{10}$: the cost of making a correct decision is smaller than the cost of making a wrong decision.)

One more important concept must now be introduced before we can discuss optimization of decision rules. We have noted that the error probabilities α and β are conditional upon either the state of nature S_0 or S_1 being true. In order to obtain an over-all probability of making an error, and more important, to obtain the average cost of a particular decision rule, it is necessary to state the a priori probabilities of S_0 and S_1 being true. In particular let $P = \text{prob. } (S_1)$, $1-P = \text{prob. } (S_0)$ be the a priori probabiliti. P then is the (perhaps subjective) probability that the target is in the cell.

In the discussion that follows we shall consider, as in the example in the previous section, the result of the search in a cell to be some observable x . The decision rule then consists of describing two exhaustive and mutually exclusive regions of the x-space X_0 and X_1 , such that, as before

$$\text{if } x \in X_0 : D_0$$
$$x \in X_1 : D_1 .$$

Again, p.d.f.s of x under the hypothesis S_0 and S_1 are assumed to be known, and are given as $p_0(x)$ and $p_1(x)$ respectively

2.4.1 Bayes Criterion

A rule that selects a strategy so that the expected cost per decision using that strategy is a minimum, compared to any other rule, is called a "Bayes" decision rule*. For a particular decision

*The naming of this criterion for the Reverend Bayes is somewhat puzzling, but the term is too familiar in the literature to change it here.

rule represented by regions X_0 and X_1, the expected cost per decision cost is seen to be given by

$$\overline{C} = \{S_0\}[C_{00}\{D_0|S_0\} + C_{10}\{D_1|S_0\}] + \{S_1\}[C_{11}\{D_1|S_1\} + C_{01}\{D_0|S_1\}] \tag{2.3}$$

where for convenience we let $\{E\} \equiv$ prob. $\{E\}$, (E is any event). Using expressions from equation (2.2) we have

$$\overline{C} = (1-P)\{C_{00}\int_{X_0} p_0(x)\,dx + C_{10}\int_{X_1} p_0(x)\,dx\} + P\{C_{11}\int_{X_1} p_1(x)\,dx + C_{01}\int_{X_0} p_1(x)\,dx\} \tag{2.4}$$

This expression is indeed the expected cost. In the first bracket in equation (2.4) the first term represents the cost of making the right decision, times the probability that this decision will be made, both conditioned on S_0 being the state of nature. To this is added the cost of making the wrong decision times the probability of making it, again conditioned upon S_0. This sum is then multiplied by the probability that S_0 will be the state of nature. The second term is a similar expectation given S_1 is true, times the probability that S_1 will be true.

The optimum Bayes decision rule is one that selects regions X_0 and X_1 such that the expression $\overline{C}$ in equation (2.4) is a minimum. It can be shown (see, for example reference (19)) that the optimum Bayes decision rule is the one that assigns X_0 and X_1 such that

$$\text{if } \frac{p_1(x)}{p_0(x)} \geq \frac{1-P}{P} \frac{C_{10}-C_{00}}{C_{01}-C_{11}} \text{ then } x \in X_1 : D_1$$

$$\text{if } \frac{p_1(x)}{p_0(x)} \leq \frac{1-P}{P} \frac{C_{10}-C_{00}}{C_{01}-C_{11}} \text{ then } x \in X_0 : D_0 . \tag{2.5}$$

The quantity $L(x) = \frac{p_1(x)}{p_0(x)}$ is known as the likelihood ratio.

In most cases of interest $L(x)$ will be monotonic in x, so that the decision rule (2.5) may be re-written

$$\text{if } x \geq x^* : D_1$$

$$x \leq x^* : D_0$$

where x^* is called the "threshold", or "bias", and is the solution to the equation

$$L(x^*) = \frac{1-P}{P} \frac{C_{10}-C_{00}}{C_{01}-C_{11}} . \tag{2.6}$$

Note that since

$$p_d = \int_{x^*}^{\infty} p_1(x)\,dx$$

$$p_f = \int_{x^*}^{\infty} p_0(x)\,dx \tag{2.7}$$

equation (2.6) may be written

$$\left. \frac{d(p_d)}{d(p_f)} \right|_{x=x^*} = \frac{1-P}{P} \frac{C_{10}-C_{00}}{C_{01}-C_{11}} . \tag{2.8}$$

Using equation (2.8) and the ROC it is quite easy to obtain the threshold, and operating point (equation 2.7) for any given device. Since the ROC is just a plot of p_d v.s. p_f, simply find the point at which the slope is equal to

$$\frac{1-P}{P}\left(\frac{C_{10}-C_{00}}{C_{01}-C_{11}}\right) .$$

This is illustrated in Figure 2.3.

The average minimum cost, using the Bayes solution then becomes

$$\overline{C}_{min} = (1-P)\{C_{00} + p_f(C_{10}-C_{00})\} + P\{C_{01} + p_d(C_{11}-C_{01})\} \qquad (2.9)$$

2.4.2 Numerical Example - The Shepherd and the Wolf

To illustrate the above development, and to provide an example for comparison of STSD with what is to follow, we shall now discuss a hypothetical search situation. An intentional attempt has been made to avoid a military example, but the reader may care to interpret the target, detection device and searcher in the example in whatever way he wishes.

The proverbial Boy who Cried Wolf comes rushing to a shepherd with the news that there is a wolf hiding near the pasture, waiting for the shepherd to go to town. The shepherd knows from past experience that the boy has probability P of telling the truth.

In order to avoid the long walk to the pasture, the shepherd has installed a microphone there, which records the pasture noises

[redacted]

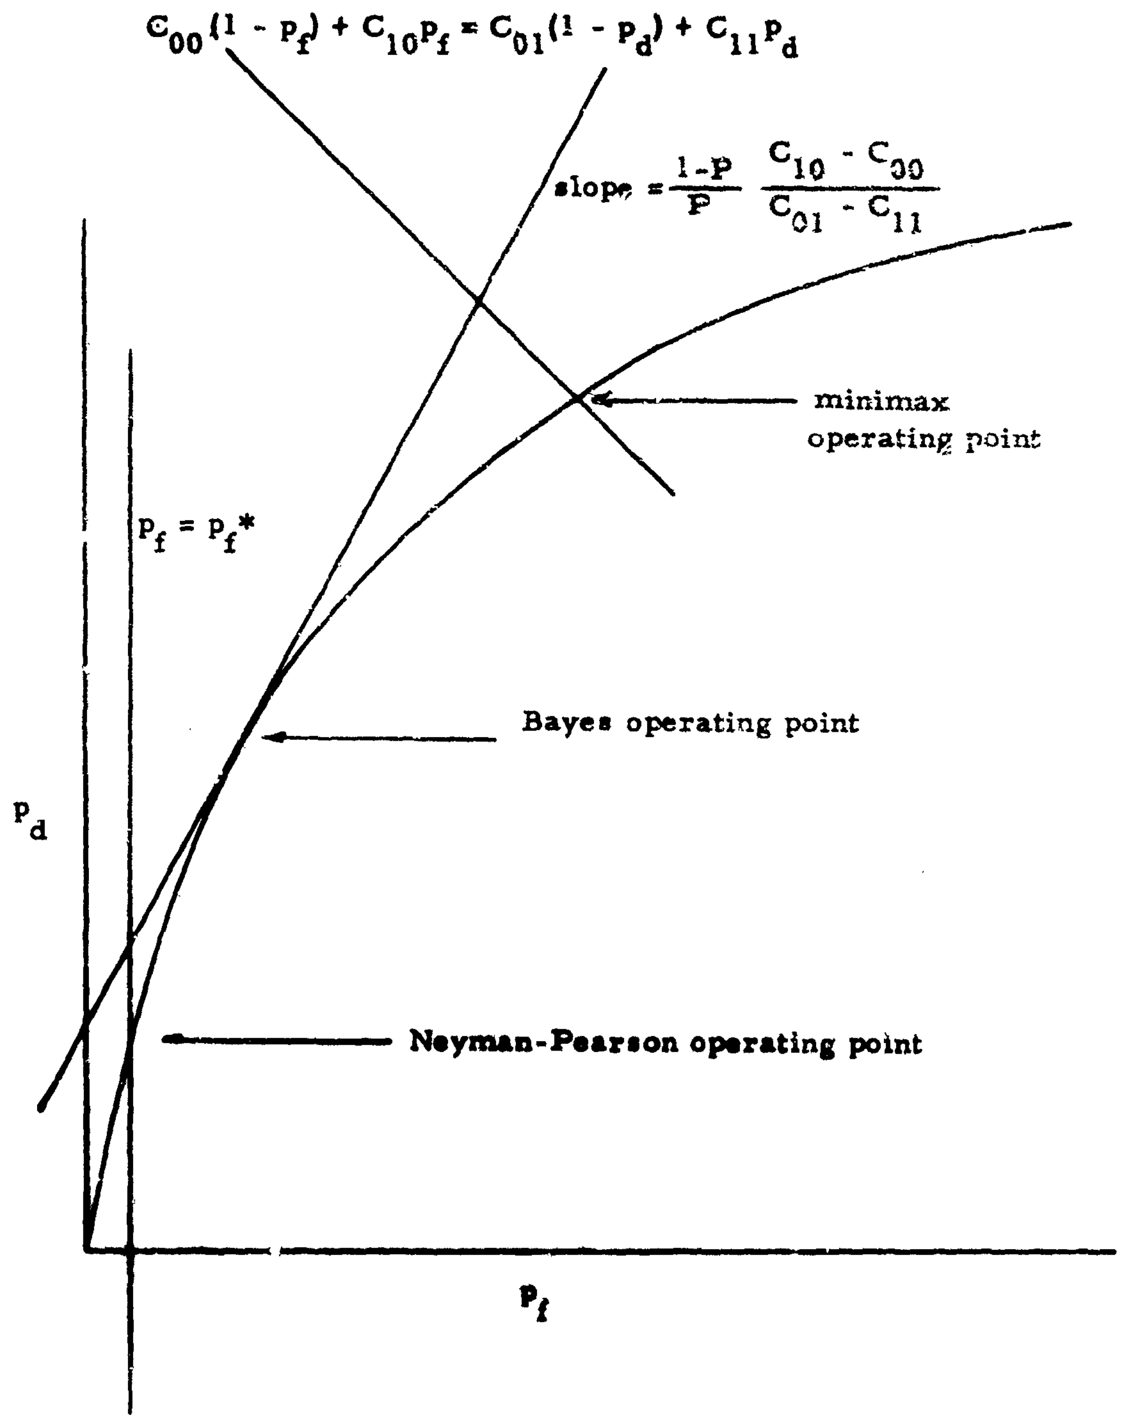

Figure 2.3 Use of the ROC to obtain Bayes, minimax and Neyman-Pearson operating points

(integrated over ten minutes) and registers the intensity on a patented "Baa-meter". It is known that if a wolf is present in the pasture, the measurement x of the meter will have the p.d.f.

$$p_1(x) = f_N(x;\ 1, 1)$$

while if there is no wolf present, the measurement x will have the p.d.f.

$$p_0(x) = f_N(x;\ 0, 1)\ .$$

Once the shepherd observes a measurement, he must decide whether to go out to the pasture and hunt the wolf, or go to town as planned. The costs of making the wrong decision are as follows:

If the shepherd goes to town he makes a \$50 sale. However, if the wolf is indeed present, a sheep worth \$100 is eaten, making a net cost of \$50.

If the shepherd goes to the pasture and the wolf is present, he kills it and collect a \$100 bounty. But whether or not the wolf is there he loses the chance to make the \$50 sale in town. Thus the cost matrix is

$$[C_{ij}] = \begin{bmatrix} -50 & 50 \\ 50 & -50 \end{bmatrix}$$

Since this cost matrix may be additively normalized, we could also write the equivalent matrix (in units of \$100):

$$[C_{ij}] = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \ - .5$$

where the matrix term is the loss due explicitly to decision errors.

By utilizing this information and substituting into equation (2.6) the shepherd finds that his optimal threshold is the solution to

$$L(x^*) = \exp(x^* - \frac{1}{2}) = \frac{1-P}{P} \quad ,$$

$$x^* = \frac{1}{2} + \ln \frac{1-P}{P} \quad .$$

If the boy has a $P=\frac{1}{2}$ of being right, then if the Baa-meter reads $x > \frac{1}{2}$, the shepherd goes to the pasture, and he goes to town if $x < \frac{1}{2}$. The resulting detection and false-alarm probabilities are

$$p_d = 1 - \operatorname{erf}(-\frac{1}{2}) = .69$$

$$p_f = 1 - \operatorname{erf}(\frac{1}{2}) = .31$$

and the average cost to the shepherd is - \$19 (a profit of \$19).

2.4.3 Other Criteria: Minimax and Neyman-Pearson

It sometimes happens that the decision maker thinks he cannot estimate the prior probabilities $1-P$ and P of the two states of nature S_0 and S_1, or that he cannot supply a cost matrix (C_{ij}), or both. We shall quickly mention here methods that have been used to overcome these difficulties. However, it is to be kept in mind that all of these methods lead to accepting <u>some</u> operating point for the system, and from the discussion in the previous section we see that this point, corresponding to a Bayes solution, implies some specific relations about costs and prior probabilities.

When the prior probability P is assumed to be unknown the "minimax" criterion calls for operating such that no matter what P is, the maximum possible loss is minimized. This argument is

usually justified by considering "nature", or the target, to be in a two-zero-sum game against the searcher, so that it tries to pick the worst possible P . A very good treatment of this approach is presented in Blackwell and Girschick (7). Because of the game characteristics, the solution may be found by finding the maximum of the minimum cost to the searcher obtained by using a Bayes' Rule as if the P was known, and then operating as if the P producing this maximum cost was in fact the real prior probability. The proper operating point may be obtained by noting that maximizing equation (2.9) with respect to P yields the expression

$$C_{00}(1-p_f) + C_{10}\,p_f = C_{01}(1-p_d) + C_{11}\,p_d$$

which is a straight line in the p_d-p_f plane. The intersection of this line with the ROC gives the minimax operating point (see Figure 2.3).

When the searcher is unwilling to supply both the prior probabilities <u>and</u> a cost matrix, then the Neyman-Pearson criterion is often used. The searcher selects some arbitrary value of false alarm probability, $p_f{}^*$, which must never be exceeded. The solution is to maximize p_d such that $p_f \le p_f{}^*$. For most simple cases, when the ROC is monotonic as in Figure 2.1, this is achieved by operating such that $p_f = p_f{}^*$, and the threshold x^* the solution to

$$\int_{x^*}^{\infty} p_0(x)\,dx = p_f{}^*$$

Other criteria are mentioned in the literature and include minimizing p_f for a fixed p_d (a sort of inverse Neyman-Pearson

procedure); assuming that $C_{01}=C_{10}$, $C_{00}=C_{11}$, and $P=\frac{1}{2}$ (the "Ideal Observer"), and so on.

In most of the work of this paper it shall be assumed that prior probabilities are available, and that specific cost models are applicable, so that the criteria mentioned above will not be useful except for possible comparative purposes It is important to note, however, that it can be shown that all of the above criteria produce a threshold decision rule. That is, the rules all require the calculation of a likelihood ratio, and comparison of it to some fixed value.

2.5 Sequential Hypothesis Testing - The sprt

The decision structures in the previous sections were all derived on the basis that the observable could be a vector $\underline{x} = (x_1, x_2, \ldots, x_n)$, so that the threshold $\underline{x}^*$ could be a n-1 dimensional figure in n-space represented by the solution to

$$\frac{p_1(\underline{x}^*)}{p_0(\underline{x}^*)} = \frac{p_1(x_1^*, x_2^*, \ldots, x_n^*)}{p_0(x_1^*, x_2^*, \ldots, x_n^*)} = \Lambda$$

where Λ is some fixed value, such as $\frac{1-P}{P}\left(\frac{C_{10}-C_{00}}{C_{01}-C_{11}}\right)$ as in section (2.4). The components of $\underline{x}$ may be n values of some measurement taken at successive times. The analysis in the previous sections holds providing that n is fixed in advance.

When the cost of taking each measurement is considered however, it may be desirable not to fix in advance the number of measurements to be taken, but to allow for the number to be determined

on the basis of what has been observed up to that time. This is called a sequential test, and the related theory has been developed primarily by Wald (47, 48). In the consideration of search problems, the theory has been specifically applied to the sequential testing of hypothesis, these usually being the two aleternatives S_0 and S_1. Again it is assumed that the target is stationary throughout the search.

The Wald sequential probability ratio test (sprt) involves not only the two possible decisions (D_0 and D_1) considered before, but also includes a third possible decision W : wait for at least one more observation.

On the basis of observing x_1 for the first measurement, the likelihood $L_1(x_1) = p_1(x_1)/p_0(x_1)$ is computed and compared to two quantities A and B $(A > B)$. The decision rule is

$$\begin{aligned} \text{if } & L_1(x_1) > A : D_1 \\ & L_1(x_1) < B :: D_0 \\ & B \leq L_1(x_1) \leq A : W . \end{aligned} \tag{2.10}$$

The D_0 and D_1 decisions are the terminal decisions, and they stop the process with a definite commitment of some sort of action. If the W decision is made, then a new measurement is taken, and a new likelihood ratio is computed. If the process has gone on for k observations, then the likelihood to be used in the next comparison is

$$L_k(x_1, x_2, \ldots, x_k) = \frac{p_1(x_1, x_2, \ldots, x_k)}{p_0(x_1, x_2, \ldots, x_k)}$$

with the decision rule the same as in (2.10), but with L_1 replaced by L_k.

The crux of the problem is calculating the values of A and B so that a given criterion is satisfied.

With the exception of Blackwell and Girshick (7), the accepted procedure has been to use a criterion based upon conditional error probabilities (α and β) only. That is, A and B are selected so that when S_0 is the state of nature the probability of reaching terminal decision D_1 is p_f, while if S_1, the probability of terminating with decision D_0 is $1-p_d$. Obtaining these required values of prob. $\{D_1|S_0\}$, and prob. $\{D_0|S_1\}$ under the sprt is in general a very difficult problem. Wald has shown, however, that the relations between A and B and p_d and p_f are governed by the inequalities

$$\left.\begin{aligned} A &\le \frac{p_d}{p_f} \\ B &\ge \frac{1-p_d}{1-p_f} \end{aligned}\right\} \qquad (2.1.$$

In order to understand the test better, let us examine the case where the observations are independent. Then we can write

$$p_i(x_1, x_2, \ldots, x_k) = \prod_{j=1}^{k} p_i(x_j)\ , \qquad i = 0, 1\ .$$

If we take logarithms of the likelihood function as well as the boundaries

$$Z_k = \ln L_k(x_1, x_2, \ldots, x_k) = \sum_{j=1}^{k} z_j$$

$$a = \ln A$$

$$b = \ln B$$

where

$$z_j = \ln [p_1(x_j)/p_0(x_j)] \tag{2.12}$$

Then the decision rule (2.10) becomes

$$\begin{aligned} \text{if } Z_k &> a \ : \ D_1 \\ Z_k &< b \ : \ D_0 \\ b \leq Z_k &\leq a \ : \ W \end{aligned}$$

where Z_k is the sum of k identically distributed random variables z_j. The density function of z_j is $g_0(\cdot)$ or $g_1(\cdot)$ as the state of nature is S_0 or S_1, where $g_i(z)\,dz = p_i(x)\,dx$ and x and z are related as in (2.12).

Thus we can view the process as a random walk of a $g_i(\cdot)$ - distributed random variable, with absorbing barriers at a and b. A further discussion of this problem is taken up in appendix A.

It can be shown that for a large class of p.d.f.s. on z the process will eventually terminate. That is

$$\text{prob.}\{b < Z_k < a \ ; \ \text{for all } k \leq n\} \rightarrow 0 \text{ as } n \rightarrow \infty \ .$$

The advantages of using the sprt is in the proven result (see references (47, 48)) that the average sample number (ASN) needed to obtain a given set of α and β error probabilities (i.e. a given p_d and p_f) is always less than the fixed sample length (say N)

required to achieve the same error probabilities, no matter what the state of nature. (Since the number of samples in the sprt is a random variable, however, it is to be noted that it will occasionally be greater than N.)

It can be shown that if the z_j are in some sense small in comparison with the distance $(a-b)$, the inequality signs in (2.11) may be replaced by equalities. With this approximation it is also possible to derive expressions for the ASN under both states of nature (see Wald (47) for the derivation):

$$\begin{aligned} ASN_0 &= [a - (1-p_f)(a-b)] \,/\, E_0(z) \\ ASN_1 &= [a - (1-p_d)(a-b)] \,/\, E_1(z) \end{aligned} \qquad (2.13)$$

where

$$E_k(z) = \int_{-\infty}^{\infty} \ln[\, p_1(x)/p_0(x)]\, p_k(x)\, dx \;; \quad k = 0, 1 \qquad (2.14)$$

is the expected value of the random variable z under S_k. When $E_0(z)$ or $E_1(z)$ are zero, these equations may be replaced by ones involving the variances.

2.5.1 Numerical Example

We return to the shepherd of section 2.4.2, who has now decided to apply the sprt, and has arbitrarily declared that the test must have $p_d \geq .69$ and $p_f \leq .31$. The random variable z of equation (2.12) becomes

$$z = \ln[\,L(x)\,] = x - \frac{1}{2}$$

so that

$$g_0(z) = f_N(z;\ -\frac{1}{2},\ 1)$$

$$g_1(z) = f_N(z;\ \frac{1}{2},\ 1)\ .$$

In addition, the assumption is made that the inequalities of equation (2.11) are equalities, so that $A = \frac{.69}{.31}$, $B = \frac{.31}{.69}$ and thus $a = .8$, $b = -.8$.

The resulting test is to compare $z = x - \frac{1}{2}$ with a and b, so that

$$\begin{aligned} \text{if } -.3 \le x \le 1.3 &: \text{take another measurement} \\ x > 1.3 &: \text{go to the pasture} \\ x < -.3 &: \text{go to town .} \end{aligned}$$

If another measurement is taken, then $z_1 + z_2 = x_1 - \frac{1}{2} + x_2 - \frac{1}{2}$ is compared to $.8$ and $-.8$, and so on.

Since $E_0(z) = -\frac{1}{2}$ and $E_1(z) = \frac{1}{2}$ the approximate formulae (2.13) give

$$ASN_0 = -[.8 - .69(1.6)]\,2 = .6$$

$$ASN_1 = [.8 - .31(1.6)]\,2 = .6$$

These results are puzzling to the shepherd, as he sees that the average number of times he will have to take measurements is less than one ! However he notices that the standard deviation of the random variable y is 1 and since $(a-b) = 1.6$, the requirement that z be

small compared to (a-b) is not fulfilled, so he does not expect equatio (2.13) to hold. The exact solution, as shown in Appendix A, is unavail

What should he do now? In fact, the choice of arbitrary p_d · p_f was really forced upon the shepherd by the constraints of the sprt. Being a practical person, he still really desires to minimize his costs. Do the measurements cost him anything? After all, if they don't, he would be satisfied with measuring all night and eventually (because of the central limit theorem) he would be certain about the state of nature Apparently, a new model and solution must be used.

2.6 Need for extension of the models

In section 2.2 it was pointed out that the defect of previous approaches to what has been conventionally called the "search" proble has been in the neglect of the possibility that the detection and sensing devices used could produce spurious signals, and hence false alarms. These theories, however, were certainly thorough in the applications of cost and utility models, as well as prior probabilities.

The non-sequential work in STSD includes false alarms, a cost structure, and cognizance of prior information, but this is limited to the simple hypothesis test. As we have seen, this contains only a rudimentary element of search, in that the decisions are restricted to terminal ones after a single observation. There is no structure, for example, that allows one to stop and decide D_0 or D_1 on the basis of no observations at all.

On the other hand, with the introduction of sprt the cost structure and prior information are lost, the decision maker is left

with the task of assigning errors of the first and second kind, and there is no consistent way to take into account the experimental cost. In addition, the sprt must be analysed in terms of the inequality (2.11), which becomes an equality only under conditions that are equivalent to requiring that $p_f \ll 1$, $1-p_d \ll 1$. Although this condition might seem desirable, we have seen in the numerical example above that minimum cost solutions do not particularly satisfy it.

Although some pioneer work by Blackwell and Girshick (7) has been done some time ago in an attempt to [illegible] "Bayes" structure on sequential hypothesis testing, the applications were not geared towards the problems of detection. The following chapters will consider a minimum cost analysis of sequential detection, under some conditions that are perhaps more reasonable than those proposed before. The technique of analysis, that of using dynamic programming, will then lead to the consideration of yet another problem, one that could not otherwise be attacked by the sprt type of approach.

Before going on to such analysis, it is of interest to mention here some work that has been published in recent years on the subject of sequential detection theory. To repeat the point made above, none of this work takes into account the overall cost of operation, and thus the results are interpreted in terms of α and β errors as well as expected test lengths.

In many realistic search problems, the target characteristics are not exactly known, and so some of the parameters in $p_1(x)$ are themselves random variables. The searcher not only must decide upon the presence or absence of such a target, but an estimate of the unknown parameters is also required. Such estimation problems

[redacted]

(sometimes referred to as "classification"), extremely difficult to analyse in a sequential manner, have been treated by Selin (41) and Turner (44).

Some work has been done in an effort to obtain exact solutions to the random walk of the sprt as discussed in section (2.5) when the p.d.f.s involved are peculiar to those found in practical detection devices. Reed (39) derives some theoretical results, and Marcum and Swerling (31), for example, produce Monte-Carlo simulations for practical cases.

Some arbitrary "many threshold" decision rules have been analysed. An example is the rule presented by Kennedy (24) which starts a sequential test only after an initial signal exceeds some fixed threshold, this latter threshold set to limit the number of times the sequential test is applied.

Bussgang (8, 9) and Middleton (34), among others, have treated sequential detection under the most general noise and signal statistics. Helstrom (20) and Preston (38) have compared the sprt to the fixed length test for practical examples of search radars.

The particularly interesting problem of sequentially testing a <u>continuous</u> signal (the experiments are not done at discrete times) has been treated by Selin (41). In this work he also indicates the optimal sequential test when the noise is correlated in time. Gray (17) has also contributed to this problem.

CHAPTER III

STATIONARY TARGET

This chapter describes a more general problem than those presented earlier, in that it provides a combination of the Distribution of Effort and STSD approaches. The mode of attack and the mathematical technique involved reflect the sequential nature of the solution as well as the minimum expected cost aspects of decision theory. The use of stochastic dynamic programming in this regard has been generally indicated by Bellman et al (4), as well as Blackwell and Girschick (7). Goode (16) has applied the principles involved, but with a model not particularly suited to search.

The strategy developed in this chapter will turn out to be a modified form of the sprt, with the cost factors and prior probabilities appearing as intrinsic parameters. In addition, the resultant minimum cost arises as a natural consequence of the calculations. This makes it possible to compare, in an efficient way, the cost of the optimal strategy with costs of certain non-optimal strategies that will be considered for practical reasons.

3.1 Problem Statement

1. The target is either present (S_1) or not present (S_0) in the cell of interest, with prior probabilities P and $1-P$ respectively, and remains so for the entire search (the target is "stationary").

2. If the cell is observed, the result is a random variable x which has p.d.f. $p_0(x)$ or $p_1(x)$ as the state of nature

is S_0 or S_1. Observations take place at unit time interval

3. After every observation (including the zeroth) the searcher makes one of the following decisions:

   D_1 = Decide target is present

   D_0 = Decide target is not present

   W = Wait for another observation.

4. The decisions D_0 and D_1 are terminal decisions, and end the process. The costs to the searcher making decision D_i given S_j is true are $C_{ij} (i, j = 0, 1)$, with $C_{11} = C_{00} = 0$

5. The decision W continues the process at least one more tim unit. The cost of this delay depends upon the state of nature, and is W_i if S_i is true $(i = 0, 1)$.

6. The objective of the searcher is to minimize the expected cost of a search. The strategy (that is, the rule for making decisions given a sequence of observations) that achieves this minimum expected cost is the "optimal" strategy.

The fact that the cost of experimentation depends upon the state of natur is what makes this model particularly applicable to search problems. For example, consider the case of active sonar search against a submarine that is suspected of being in a missile-firing position. Clearly, the searcher would like to make a decision as soon as possible if the submarine is indeed present, whereas if the indications are that the submarine is not present, the searcher could afford to spend more time making sure. Thus in many search situations we should expect to have $W_1 > W_0$.

*This assumption is the equivalent of zero-normalizing the cost for any given P, and is made to simplify the algebra.

We shall now write a functional equation from which the optimal strategy car be obtained, the equation being a straight-forward application of Bellman's Principle of Optimality (3).

Let $C(P)$ be the minimum cost of search obtained by following the optimal strategy, where P is the a priori probability that the target is present. This minimum cost will be the smaller of the costs resulting from the three possible decisions that can be made at that time:

If the decisions are D_0 or D_1, the costs are due to terminal wrong decisions: PC_{01} and $(1-P)C_{10}$ respectively.

If the decision is W, the cost is $PW_1 + (1-P)W_0$ plus the cost of continuing from then on, having observed some value of x. bability of observing a value between x and $x+dx$ is $Pp_1(x)\,dx+(1-P)\,p_0(x)$. Having observed the value x, however, the probability now of the target being present is

$$\{S_1|x\} = \frac{\{x|S_1\}\{S_1\}}{\{x|S_1\}\{S_1\} + \{x|S_0\}\{S_0\}} = \frac{p_1(x)P}{p_1(x)P + p_0(x)(1-P)} .$$

The equation for $C(P)$ may thus be written

$$C(P) = \min\begin{cases} (1-P)C_{10} & : D_1 \\ PC_{01} & : D_0 \\ PW_1+(1-P)W_0+\int_{-\infty}^{\infty}[Pp_1(x)+(1-P)p_0(x)]\, C\left[\dfrac{p_1(x)P}{p_1(x)P+p_0(x)(1-P)}\right]dx & : W \end{cases}$$

For notational convenience let us define the minimum cost of terminal decision to be $T(P)$

$$T(P) = \min \begin{cases} (1-P)\, C_{10} & : D_1 \\ \\ P\, C_{01} & : D_0 \end{cases}$$

and the unconditional p.d.f. of receiving a value x to be $g(x)$

$$g(x) = P\, p_1(x) + (1-P)\, p_0(x)$$

so that equation (3.1) may be re-written

$$C(P) = \min \begin{cases} T(P) \\ \\ P\, W_1 + (1-P)\, W_0 + \displaystyle\int_{-\infty}^{\infty} g(x)\, C\left[\frac{p_1(x)\, P}{g(x)}\right] dx\,. \end{cases} \tag{3.2}$$

3.2 Arbitrary Truncation of the Search

In order to find a way to solve equation (3.2) we shall consider the following arbitrary truncation of the search: At the start of the process, the searcher is told that he has only n available possible observations (or equivalently, time units) left. If the decision W is made, at the next decision there will be only $n-1$ possible observations left, and so on. If $n=0$, then one of the terminal decisions must be made.

If we define $C_n(P)$ to be the minimum cost of search given that there are n available observations remaining before a terminal decision must be made, we may write the equivalent of equation (3.2) as

$$C_n(P) = \min \begin{cases} T(P) & : D_0 \text{ or } D_1 \\ PW_1+(1-P)W_0+\int_{-\infty}^{\infty} g(x)\, C_{n-1}\left[\frac{p_1(x)\,P}{g(x)}\right] dx & : W \end{cases} \tag{3.3}$$

The boundary condition on this equation is given when $n=0$, for then a terminal decision is required and so

$$C_0(P) = T(P) = \min \begin{cases} (1-P)\,C_{10} \\ P\,C_{01} \end{cases} = \begin{cases} (1-P)\,C_{10} & P \geq P^* : D_1 \\ P\,C_{01} & P \leq P^* : D_0 \end{cases} \tag{3.4}$$

where

$$P^* = \frac{C_{10}}{C_{10} + C_{01}} \quad .$$

It is of interest to go through the calculation of $C_1(P)$ in some detail, since the behavior of the terms affords an insight to the decision structure of the solution, and to the form of the minimum cost as a function of P.

In order to calculate $C_1(P)$ we must first evaluate $C_0\left[\frac{p_1(x)\,P}{g(x)}\right]$. This latter expression is the minimum cost given that one observation x has been taken and that a terminal decision must be made. From equation (3.4), we find, after some manipulation,

$$C_0\left[\frac{p_1(x)P}{g(x)}\right]=\begin{cases}\dfrac{(1-P)\,p_0(x)\,C_{10}}{g(x)} & \dfrac{p_1(x)}{p_0(x)}\geq\dfrac{1-P}{P}\,\dfrac{P^*}{1-P^*}\ :\ D_1\\[2ex] \dfrac{P\,p_1(x)\,C_{01}}{g(x)} & \dfrac{p_1(x)}{p_0(x)}\leq\dfrac{1-P}{P}\,\dfrac{P^*}{1-P^*}\ :\ D_0\end{cases}\tag{3.5}$$

If we let $\frac{1-P}{P}\,\frac{P^*}{1-P^*}=\Lambda(P)$, and recall that $p_1(x)/p_0(x)=L(x)$ the likelihood ratio, we see that the rule presented by equation (3.5) is identical with equation (2.5). (We see now that this corresponds to the solution for the truncated search, and given that an observation was taken: two restrictions that are not contained in the general problem statement.) In particular, by letting x^* again be the solution to

$$L(x^*)=\Lambda(P)$$

and putting the results of equation (3.5) into equation (3.3) (with $n=1$), we obtain

$$C_1(P)=\min\begin{cases}T(P)\\[2ex] PW_1+(1-P)W_0+\displaystyle\int_{-\infty}^{x^*}PC_{01}\,p_1(x)\,dx+\int_{x^*}^{\infty}(1-P)\,p_0(x)\,C_{10}\,dx\end{cases}\ .\tag{3.6}$$

The relation between this search and the one described in section 2.4 can now be further explored. If we let $W_1=W_0=0$, (and recall that we have let $C_{00}=C_{11}=0$) and neglect the possibility

of a terminal decision before taking any observations (the $T(P)$ term) then equation (3.6) is identical to equation (2.9). If we include these factors however, as required by the general problem, the minimum costs and strategies are quite different.

The decision structure presented by equation (3.6) may be easily visualized by drawing a sketch of the terms in the right hand side, as in Figure 3.1. $T(P)$ is seen to be two straight lines meeting at $P=P^*$. Let us define $G(P)$ to be the bottom of the right hand side of equation (3.3). Then $G_1(P)$, the lower part of equation (3.6) can be seen to have the properties as shown in Figure 3.1.

$$G_1(0) = W_0$$

$$G_1(1) = W_1 \quad .$$

The interception of $T(P)$ and $G_1(P)$ occurs at two points (if at all), these two points being at $P = \gamma_1$ and $P = \delta_1$ $(\gamma_1 \le \delta_1)$.

From the figure we can see that $T(P) < G_1(P)$ in the two regions $0 \le P < \gamma_1$ and $\delta_1 < P \le 1$, so that

$$C_1(P) = \begin{cases} (1-P)C_{10} & \delta_1 \le P \le 1 \ : \ D_1 \\ P C_{01} & 0 \le P \le \gamma_1 \ : \ D_0 \\ P W_1 + (1-P) W_0 + P C_{01} \displaystyle\int_{-\infty}^{x^*} p_1(x)\,dx + (1-P) C_{10} \int_{x^*}^{\infty} p_0(x)\,dx & \gamma_1 \le P \le \delta_1 \ : \ W \end{cases} \tag{3.7}$$

where δ_1 is the solution to

$$(1-\delta_1) C_{10} = G_1(\delta_1) \tag{3.8}$$

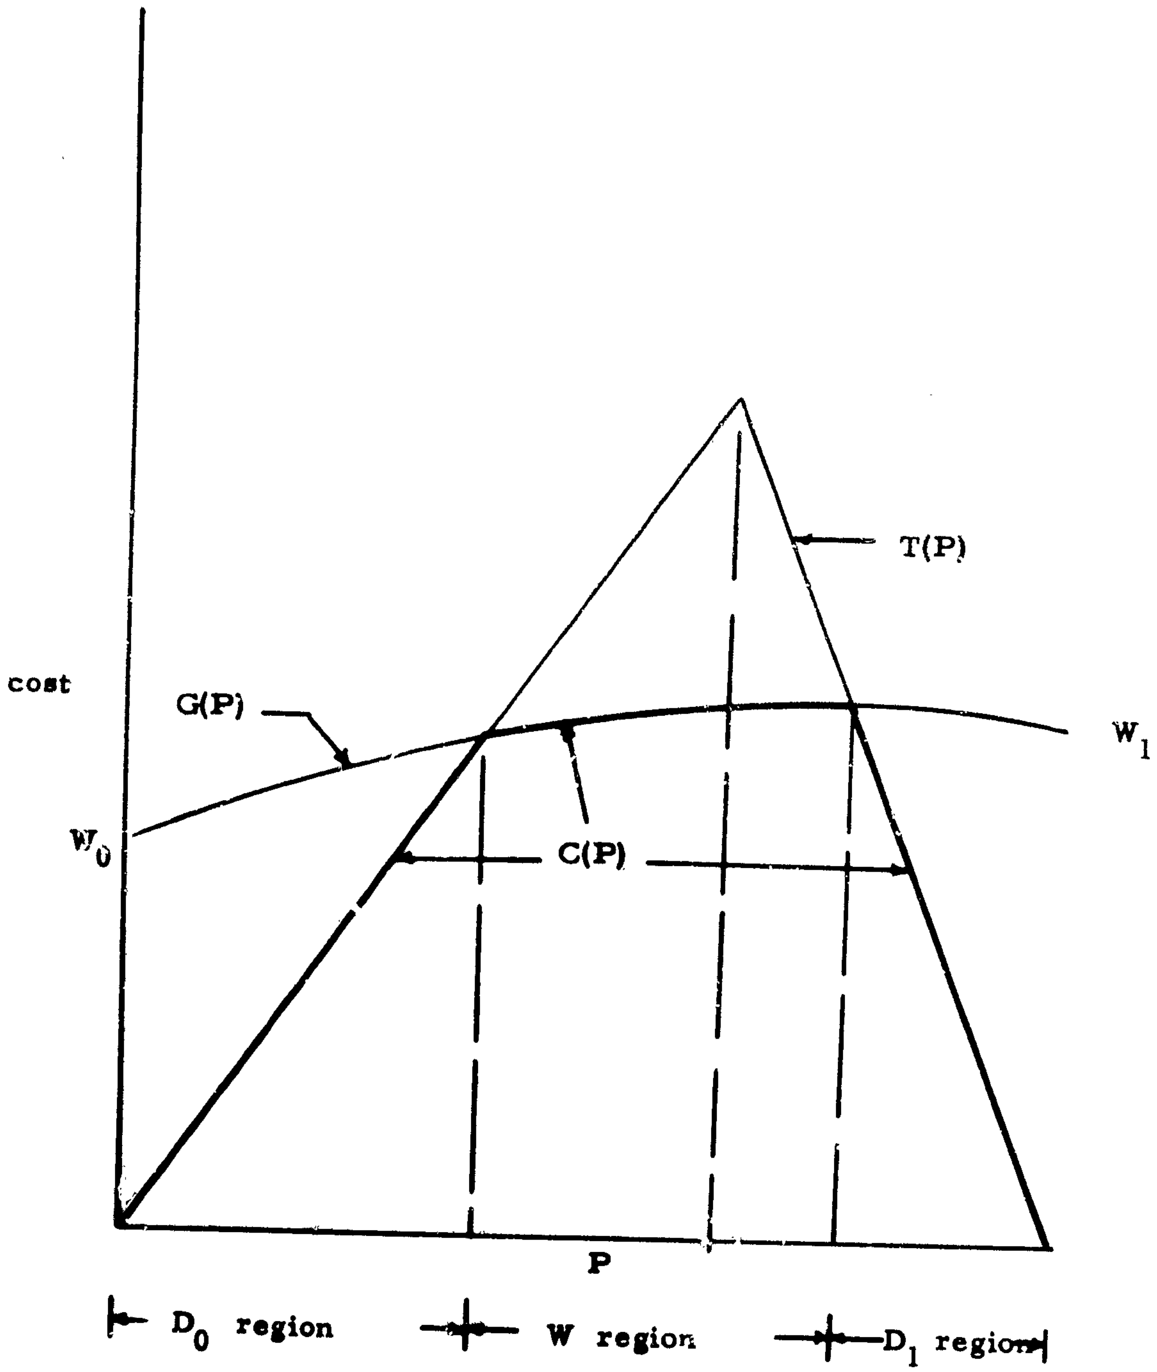

Figure 3.1 The construction of C(P)

and γ_1 is the solution to

$$(1-\gamma_1) C_{10} = G_1(\gamma_1) . \tag{3.9}$$

Since x^* is a function of P, and the integrals in $G_1(\cdot)$ are usually not expressable in closed form, the solutions of equations (3.7) and (3.8) must be obtained numerically.

3.3 C(P) as the Limit of the Truncated Search

Now that $C_1(P)$ has been obtained, we can proceed to iteratively solve equation (3.3) by letting $n = 2, 3, \ldots$ and so on. These iterations will also produce the decision regions bounded by the interception of $T(P)$ with $G_2(P), G_3(P), \ldots$ at points γ_2, δ_2 ; γ_3, δ_3 ; etc.

In particular, we may re-write equation (3.3)

$$C_n(P) = \begin{cases} (1-P) C_{10} & \delta_n \leq P \leq 1 \\ P C_{01} & 0 \leq P \leq \gamma_n \\ P W_1 + (1-P) W_0 + \int_{-\infty}^{\infty} g(x) C_{n-1}\left(\frac{p_1(x) P}{g(x)}\right) dx & \gamma_n \leq P \leq \delta_n \end{cases}$$

with the boundary condition

$$C_0(P) = \begin{cases} (1-P) C_{10} & P^* \leq P \leq 1 \\ \\ P C_{01} & 0 \leq P \leq P^* \end{cases}$$

where

$$P^* = \delta_0 = \gamma_0 = \frac{C_{10}}{C_{10} + C_{01}}$$

If we postulate that in the limit as $n \to \infty$ $C_n(P)$ approaches some function of P independent of n, then by letting $n \to \infty$ in both sides of equation (3.3), we see that this function is defined by equation (3.2) to be $C(P)$. Thus, if in fact this limit exists, $C(P)$ can be calculated to whatever degree of accuracy that is desired by successive iterations of equation (3.3).

What is more important, from the point of view of the strategy that is associated with $C(P)$, is the fact that if $C(P)$ can be obtained to any degree desired, then the points $\gamma = \lim_{n \to \infty} \gamma_n$ and $\delta = \lim_{n \to \infty} \delta_n$ defining the decision regions can also be obtained as closely as desired.

It remains, then, to prove that successive iterations of equation (3.3) will produce a $C_n(P)$ that will converge to some $C(P)$, and that the decision points γ and δ do exist as limits of γ_n and δ_n, and in a non-degenerate way (that is, $\gamma > 0$, $\delta < 1$).

3.3.1 Convergence Proofs for $C_n(P)$, γ_n and δ_n

We consider the defining equation of $C_n(P)$

$$C_n(P) = \min\left[T(P), G_n(P)\right] \qquad 0 \le P \le 1$$

where

$$T(P) = \min\left[(1-P)C_{10}, PC_{01}\right]$$

$$G_n(P) = PW_1 + (1-P)W_0 + \int_{-\infty}^{\infty} g(x)\, C_{n-1}\left[\frac{p_1(x)P}{g(x)}\right] dx ,$$

and the points γ_n and δ_n are defined by

$$\gamma_n C_{01} = G_n(\gamma_n)$$

$$(1-\delta_n)C_{10} = G_n(\delta_n) .$$

The costs C_{01}, C_{10}, W_1 and W_0 are all non-negative, and $g(\cdot)$ and $p_1(\cdot)$ are p.d.f.s.

<u>Theorem A</u>: $C_n(P) \geq 0$ for all $n \geq 0$

(This theorem provides a lower bound for the iteration process.)

<u>Proof</u>

1. Since C_{10}, $C_{01} \geq 0$, then $T(P) \geq 0$,
2. $C_0(P) = T(P)$ by definition so that $C_0(P) \geq 0$,
3. If $C_n(P) \geq 0$ for all P, then any average of $C_n(P)$ over P is ≥ 0,
4. $G_{n+1}(P)$ is the sum of $P W_1 + (1-P) W_0$ plus an average of $C_n(\cdot)$, all ≥ 0, so that $G_{n+1}(P) \geq 0$.
5. Finally $C_{n+1} = \min[T(P), G_{n+1}(P)] \geq 0$ by (1) and (4), and the induction is complete starting with the result of (2).

<u>Theorem B</u>: $C_{n+1}(P) \leq C_n(P)$

(This theorem, coupled with Theorem A, establishes the absolute convergence of $C_n(P)$ to some limit function (defined as $C(P)$).)

<u>Proof</u>

1. $C_0(P) = T(P)$ by definition,
2. $C_1(P) = \min[T(P)] = \min[C_0(P), G_1(P)] \leq C_0(P)$,
3. Suppose $C_{n+1}(P) - C_n(P) \leq 0$ for some n,

4. $G_{n+2}(P) - G_{n+1}(P) = \int_{-\infty}^{\infty} g(x)\, C_{n+1}\left[\frac{p_1(x)\,P}{g(x)}\right] dx - \int_{-\infty}^{\infty} g(x)\, C_n\left[\frac{p_1(x)}{g(x)}\right.$

$$= \int_{-\infty}^{\infty} g(x) \left\{ C_{n+1}\left[\frac{p_1(x)\,P}{g(x)}\right] - C_n\left[\frac{p_1(x)\,P}{g(x)}\right] \right\} d$$

by (3) and the fact that $g(x) \geq 0$ for all x. So we conclude that $G_{n+2}(P) \leq G_{n+1}(P)$.

5. Finally, $C_{n+2}(P) = \min[T(P), G_{n+2}] \leq \min[T(P), G_{n+1}(P)$ $C_{n+1}(P)$ by (4), and the induction is complete starting with (2.

Theorem C: a) $\gamma_{n+1} \leq \gamma_n$

b) $\delta_{n+1} \geq \delta_n$

(This theorem, in conjunction with Theorem D, establishes the convergence of the decision points γ_n and δ_n to their respective limits γ and δ.)

Proof

(The proof is given for part b), the proof of a) being essentially the same)

1. $G_n(\delta_{n+1}) \geq G_{n+1}(\delta_{n+1})$ by (4) of theorem B.
2. $C_{10}(1-\delta_{n+1}) = G_{n+1}(\delta_{n+1})$ by definition.
3. $G_n(P) < (1-P)\,C_{10}$ for $P < \delta_n$ by the defining equation.
4. To prove by contradiction, we suppose $\delta_n > \delta_{n+1}$.
5. By (3), (2) and (1) we have the contradiction
$G_n(\delta_{n+1}) < (1-\delta_{n+1})\,C_{10} = G_{n+1}(\delta_{n+1}) \leq G_n(\delta_{n+1})$.

__Theorem D__: If $W_1 > 0$, $W_0 > 0$ then

a) $\gamma_n \geq \Gamma > 0$

b) $\delta_n \leq \Delta < 1$ for all n

where Γ and Δ are explicit functions of the cost terms. (This theorem shows that the decision regions described by the decision points γ_n and δ_n are bounded by terms greater than 0 and less than 1 respectively, and so that for sufficiently large or small P a terminal decision will always be appropriate. Since the theorem holds for all n, it holds for $\gamma = \lim_{n\to\infty} \gamma_n$ and $\delta = \lim_{n\to\infty} \delta_n$.)

__Proof__

1. $G_n(P) = P W_1 + (1-P) W_0 + \int_{-\infty}^{\infty} g(x)\, C_{n-1}\left[\frac{p_1(x) P}{g(x)}\right] dx \geq P W_1 + (1-P) W_0$.

2. Since δ_n is the solution of $(1-\delta_n) C_{10} = G_n(\delta_n)$ we have $(1-\delta_n) C_{10} \geq \delta_n W_1 + (1-\delta_n) W_0$ so that

$$\delta_n \leq \frac{C_{10} - W_0}{C_{10} - W_0 + W_1} \equiv \Delta \quad .$$

3. Since $\Delta < 1$ for $W_1 > 0$, b) is proven.

4. γ_n is the solution of $\gamma_n C_{01} = G_n(\gamma_n)$ so

$$\gamma_n C_{01} \geq \gamma_n W_1 + (1-\gamma_n) W_0$$

$$\gamma_n \geq \frac{W_0}{C_{01} + W_0 - W_1} \equiv \Gamma \quad .$$

5. Since $\Gamma > 0$ for $W_0 > 0$, a) is proven.

Note: When $W_1 = 0$ or $W_0 = 0$ the decision regions are semi-degenerate, i.e. $\Delta \leq 1$ or $\Gamma \geq 0$. This is discussed in appendix B.

Theorem E: $C_n(P) > 0$ for at least some P

(This theorem is immediately proven by noting that $C_n(\Gamma) = \Gamma C_{01} > 0$ for all n. The result provides the non-degeneracy of $C_n(P)$, and in particular the non-degeneracy of $C(P)$.)

3.4 Implementation of the Optimal Search

Before going on to a specific example of the calculation of $C(P)$ and the associated strategy, it is instructive to consider the implementation of the resulting strategy. We assume, then, that $C(P)$, γ and δ have been obtained. The search process proceeds as follows.

1. If $0 \leq P < \gamma$ or $\delta < P \leq 1$, make the appropriate terminal decision. The search is over, and the expected cost is $T(P)$.

2. If $\gamma \leq P \leq \delta$, then take an observation x_1. Calculate the a posteriori probability that the target is present:

$$\{S_1 | x_1\} = \frac{p_1(x_1)P}{g(x_1)} = \frac{p_1(x_1)P}{p_1(x_1)P + p_0(x_1)(1-P)} .$$

Compare this new probability with γ and δ. This is the equivalent of comparing the likelihood ratio $L(x_1) = p_1(x_1)/p_0(x_1)$ with the values $\frac{1-P}{P}\frac{\gamma}{1-\gamma}$ and $\frac{1-P}{P}\frac{\delta}{1-\delta}$ with the result

$$\text{if } L(x_1) < \frac{1-P}{P}\frac{\gamma}{1-\gamma} \qquad : D_0$$

$$L(x_1) > \frac{1-P}{P}\frac{\delta}{1-\delta} \qquad : D_1$$

$$\frac{1-P}{P}\frac{\gamma}{1-\gamma} \leq L(x_1) \leq \frac{1-P}{P}\frac{\delta}{1-\delta} \qquad : W$$

3. If $L(x_1)$ lies in the W region, take another observation x_2, and compare

$$L(x_1, x_2) = \frac{p_1(x_1)\, p_1(x_2)}{p_0(x_1)\, p_0(x_2)}$$

with the values $\frac{1-P}{P} \frac{\gamma}{1-\gamma}$ and $\frac{1-P}{P} \frac{\delta}{1-\delta}$ and so on.

We see, then, that we have generated a Wald sprt equivalent to the one described by equation (2.10), but with the decision points A and B given by

$$A = \frac{1-P}{P} \frac{\delta}{1-\delta}$$

$$B = \frac{1-P}{P} \frac{\gamma}{1-\gamma}$$

The advantages of the search just derived over the sprt are apparent in three important respects. First, the values of γ and δ, and thus the decision thresholds, depend by definition upon the various cost factors and a priori probabilities involved in the search. The previous method of assigning arbitrary α and β errors to determine A and B can be now examined for consistency, if not completely replaced. Second, if it is at all necessary to truncate the search, the decision points described by γ_n and δ_n are the result of a straightforward and well defined optimization process, whereas the truncation of Wald-type tests in the literature are rather arbitrary.

The third advantage of the approach just developed is the important fact that not only is the strategy calculated, but $C(P)$, the minimum cost obtained by using that strategy, is a natural by-product. Thus, if we wish to compare two different systems, it is only necessary to compare their respective $C(P)$'s, and choose

the one which has the higher $C(P)$ at the P of interest. This point will be the subject of some later sections, in which certain near-optimal strategies are considered and compared.

3.5 Examples and Sample Calculations

With the assurance of section (3.3) that the process will converge, it is extremely easy to successively iterate equation (3.3) to any desired degree of accuracy. In particular, by referring to the proof of theorem B, we see that if $C_n(P)$ is at most ϵ away from $C_{n-1}(P)$, then $G_{n+1}(P)$ will be within ϵ of $G_n(P)$, forcing a like bound in the maximum difference between $C_{n+1}(P)$ and $C_n(P)$. Thus, as in the work that follows, we may select some ϵ and stop the iteration when $\max_P \{C_n(P) - C_{n+1}(P)\} < \epsilon$.

Another helpful computational technique is available. This involves approaching the limit $C(P)$ from below rather than from above by starting out with $C_0(P) \equiv 0$ rather than $C_0(P) \equiv T(P)$ as before. That this process also converges, and to the same $C(P)$, is easily shown by simple modifications to the proofs of theorems A and B. This "convergence from below" suffers in not having any intuitive basis for the index of the convergence, but has the advantage of converging more rapidly than the standard way for very low-minimum cost systems. It also serves to check the accuracy of whatever numeric approximations might be made in the calculations, by reaching $C(P)$ independently.

3.5.1 Known Signal in Gaussian Noise

For general illustrative purposes, we shall consider the simple example of the search for a target having known characteristics

in the presence of additive noise having known (statistical) properties. In particular, let the known target signal be $s(t)$, $0 \leq t \leq T$, with

$$S = \int_0^T s^2(t)\, dt$$

the signal energy. We assume the noise is Gaussian in magnitude with auto-covariance $\frac{N}{2}\,\delta(t)$, where $\frac{N}{2}$ is the known noise energy density. It can then be shown (Helstrom (19), for example) that if a statistic x is taken to be proportional to the cross-correlation of $s(t)$ with the received signal $v(t)$, i.e.

$$x = \sqrt{\frac{2}{NS}} \int_0^T s(t)\, v(t)\, dt$$

then x is sufficient and has the p.d.f.s. under S_0 and S_1 of

$$\begin{aligned} p_0(x) &= f_N(x;\ 0, 1) \\ p_1(x) &= f_N(x;\ \mu, 1)\ , \end{aligned} \tag{3.11}$$

and

$$L(x) = \frac{p_1(x)}{p_0(x)} = \exp\left(\mu x - \frac{\mu^2}{2}\right)$$

where

$$\mu = \sqrt{\frac{2S}{N}}\ .$$

The literature contains many analyses based upon this model of signal and noise characteristics. It has the advantage of being a good approximation to many realistic situations. In addition,

the "strength" of a signal, relative to noise may be represented the single variable μ , the familiar signal-to-noise ratio. As w have seen in the discussion of the ROC in Chapter II, an increase μ (all other system variables remaining constant) is always de… but μ is often fixed for any given detection device and target.

The convergence of $C_n(P)$ to $C(P)$ from above an… is illustrated in Figure 3.2, with the parameter values $C_{01} = C_…$ $W_0 = W_1 = 1$ and $\mu = 1$. Table 3.1 shows the convergence of and δ_n for the same parameter values. Figure 3.3 shows … for various other values of W_0 , W_1 and μ .

The calculations were performed on an IBM 7090 at th… M.I.T. Computation Center, and consisted of simple iterations … equation (3.3). For the examples in Figure 3.2 and 3.3, the P- was represented by points at intervals of .01, and $C_{n-1}(P)$ was linearly interpolated between them. The integration was done nu… rically by simple trapesodial addition to an error of roughly .1%. The value of ϵ used was .001, and as can be seen from the figur the convergence is quite rapid. A flow chart of the computer pro involved appears in appendix C.

3.5.2 Numerical Example

The shepherd of the previous examples now realizes th… his experimentation (observation of readings on the baa-meter) dc cost something. In fact, because of the rural location, electricity is very expensive, and he estimates that the meter operating cost is \$10 for every ten-minute integration period (a fixed time unit

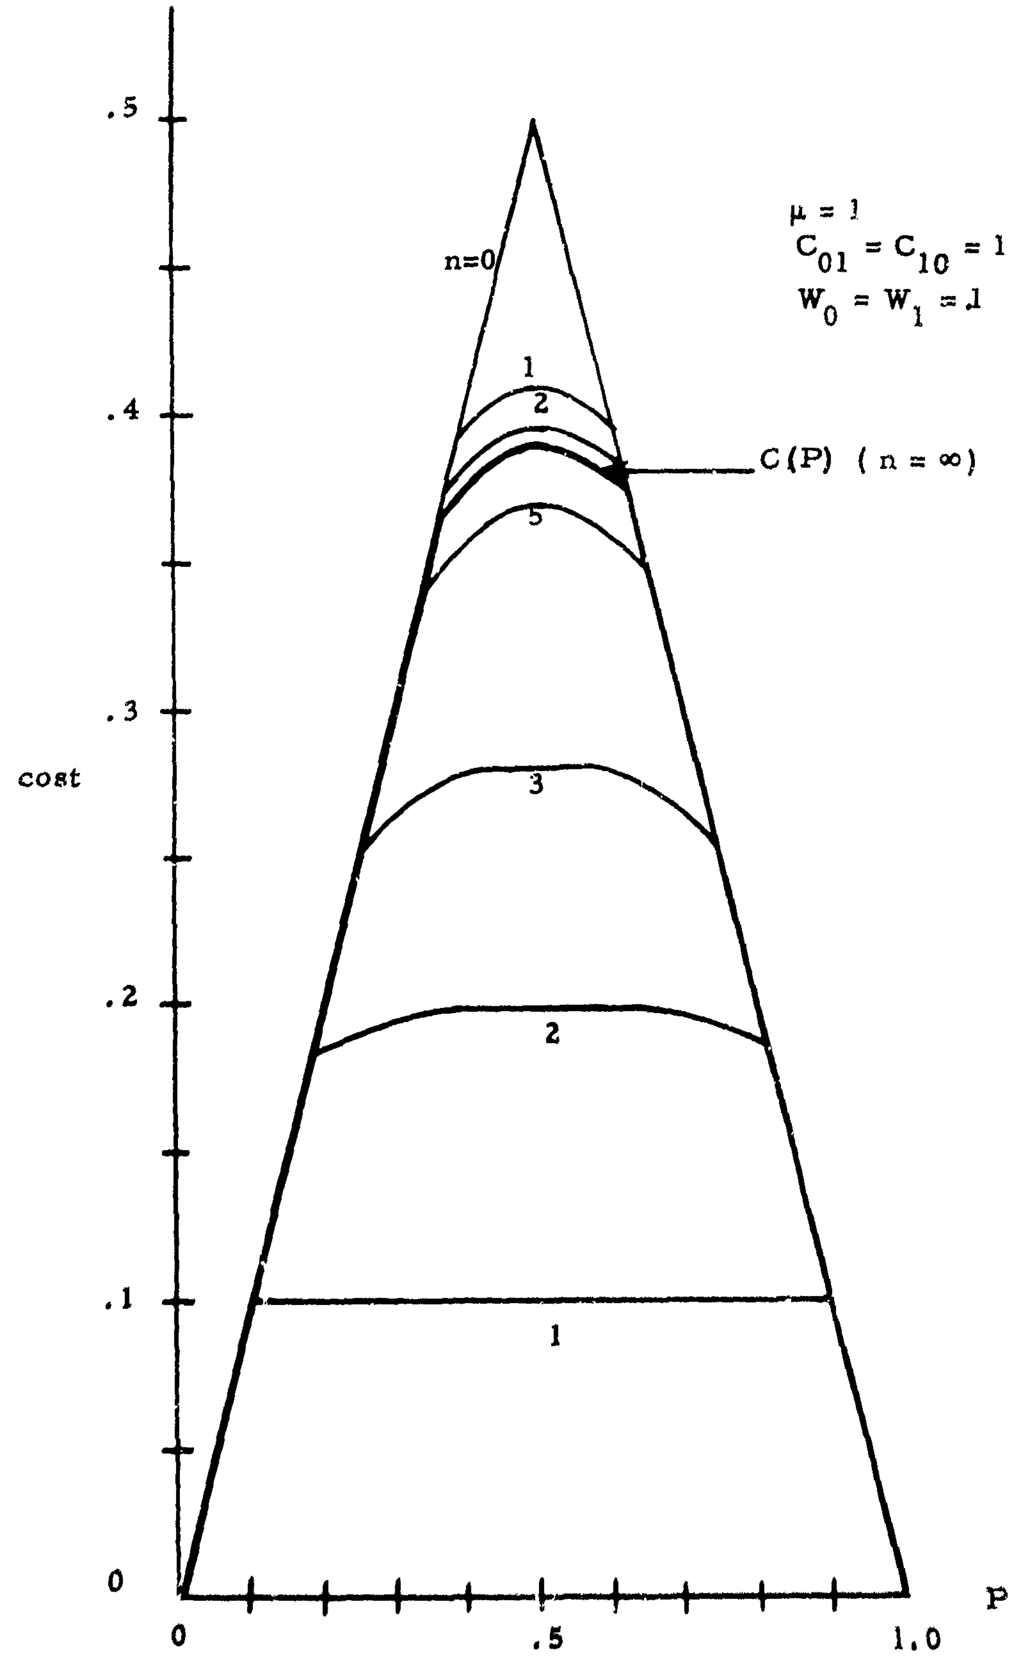

Figure 3.2 Convergence of $C_n(P)$ from above and below

Iteration Number	$C_0(P) = T(P)$		$C_0(P) = 0$	
n	γ_n	δ_n	γ_n	δ_n
0	.500	.500	--	--
1	.389	.606	.100	.900
2	.373	.621	.183	.814
3	.369	.627	.250	.744
4	.366	.629	.304	.689
5	.365	.629	.339	.655
6			.354	.639
7			.361	.632
8			.364	.631
9			.365	.630

TABLE 3.1

Convergence of γ_n and δ_n , $\mu = 1$, $W_0 = W_1 = 1$, $C_{01} = C_{10} = 1$.

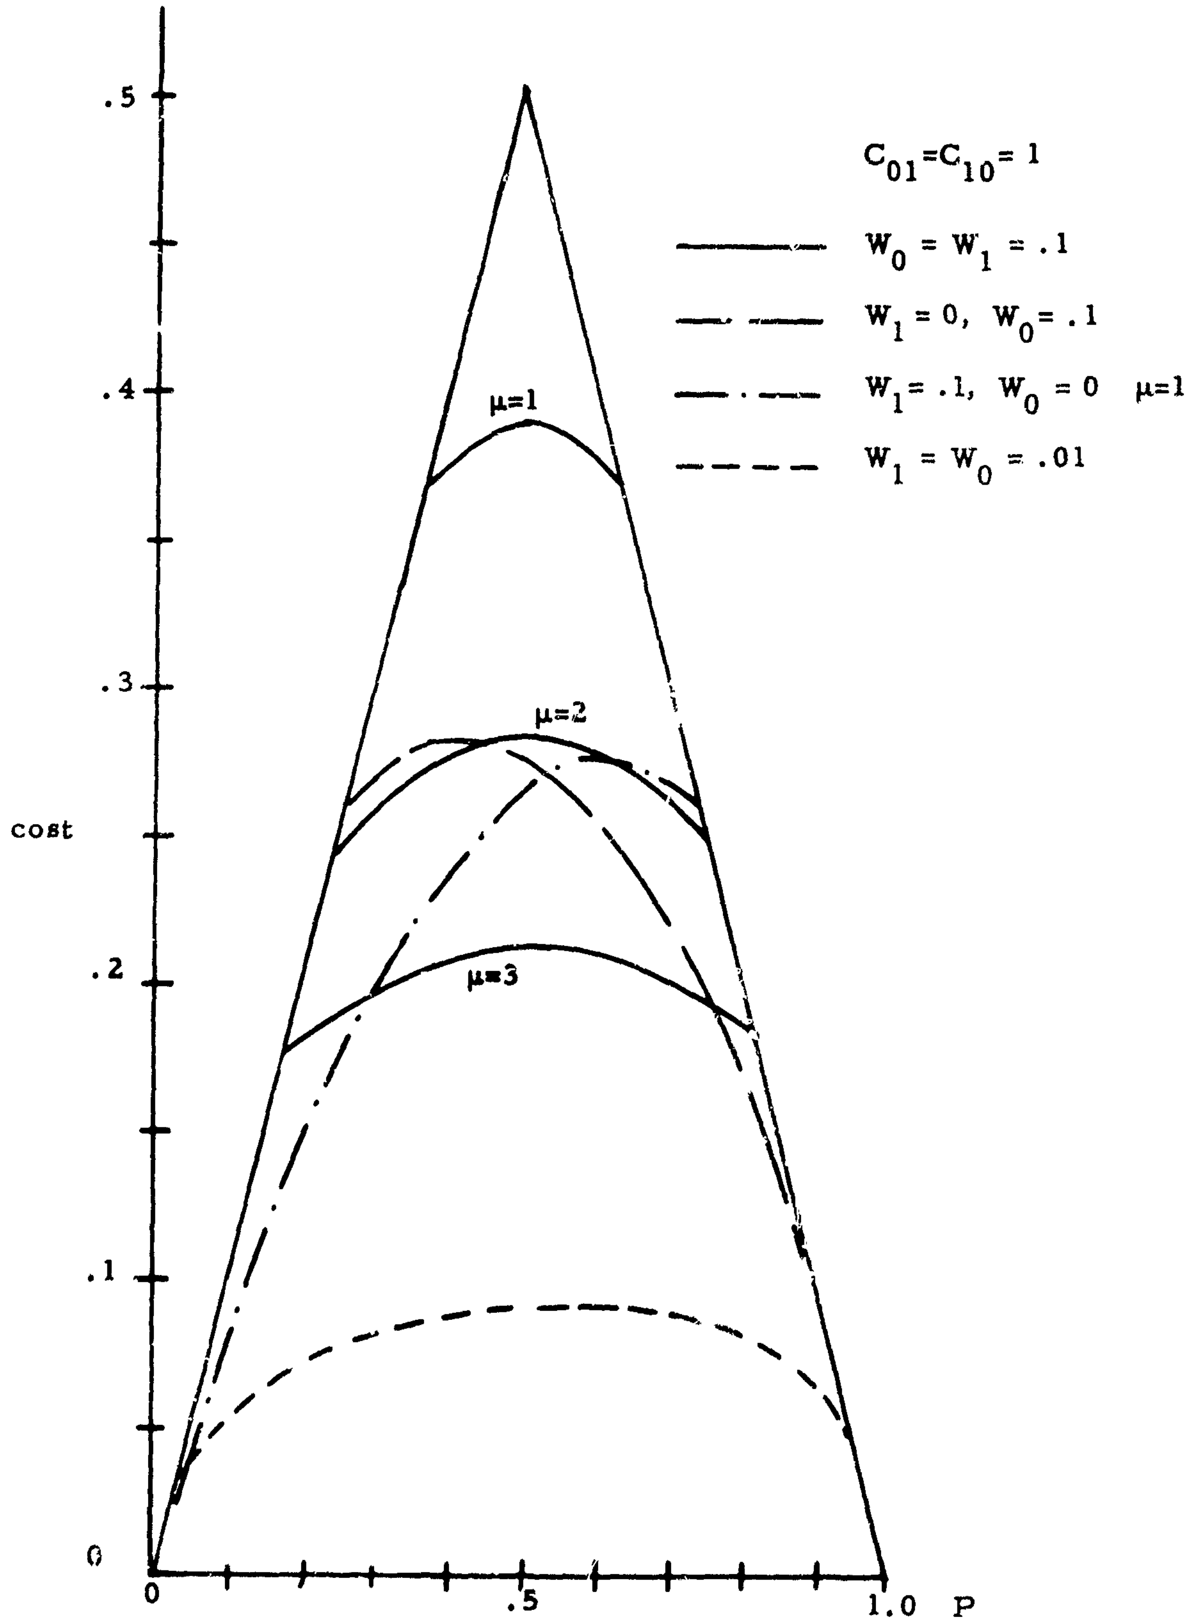

Figure 3.3 Example of C(P)

required by the meter for each observation). Using the reduced cost matrix of section (2.4.2) (in units of $100)

$$[C_{ij}] = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} -.5$$

he notes that $C_{01} = C_{10} = 1$, and the meter operating cost assigns $W_0 = W_1 = .1$. Figure 3.2 and Table 3.1 give the strategy:

1. If $P > .63$, don't take any measurements at all, and go directly to the pasture

2. If $P < .37$, don't take any measurements at all, and go directly to town

3. If $.37 \leq P \leq .63$, take a measurement x_1 . Then compare $L(x_1) = \exp(x_1 -.5)$ with the value $\frac{1-P}{P}\frac{\gamma}{1-\gamma} = \frac{1-P}{P}(.59)$, and $\frac{1-P}{P}\frac{\delta}{1-\delta} = \frac{1-P}{P}(1.85)$

4. If $\exp[x_1 -.5]$ is outside the range $(\frac{1-P}{P}(.59), \frac{1-P}{P}(1.85))$, take the appropriate terminal decision. If $\exp[x_1 -.5]$ is within that range, take another observation x_2 , and compare $\exp[(x_1 -.5)+(x_2 -.5)]$ with that range, and so on.

Thus, if $P = \frac{1}{2}$, the optimal decision rule requires at least one observation, and offers an expected loss (in dollars) due to wrong decision of 39 (from Figure 3.2), or a total loss of -50+39=-11 . By using the STSD decision rule as in section (2.4.2), the cost was -19 , to which we must add 10 , the cost of the required experiment for a total of -9 . The optimal sequential rule thus saved $2 .

The saving becomes even more substantial as P approaches 0 or 1, and is \$10 (the cost of the needless experiment required by STSD) for the certainty prior information conditions $P = 0$ or 1. This is shown in Figure 3.4.

If the cost of experimentation were much smaller (say \$1 per observation, so that $W_0 = W_1 = .01$), then the advantage of using the optimal sequential search becomes even more pronounced, as is also seen in Figure 3.4

3.6 Fixed Threshold Sequential Search

Now we shall consider some practical limitations to the applicability of the optimal search just derived. We have assumed that the result of any observation is the measurement of some continuous random variable, the value observed being called x. In order to implement the search, it is necessary to be able to calculate $L(x)$ in order to obtain the posteriori probability needed for comparison to γ and δ.

Many physical detection systems, however, are limited in their capability to measure x as a continuous variable. Others have very limited mathematical capability (whether space-cost limited computers, or time-ability limited humans) and the storage and calculation of continuous likelihood functions are beyond them.

A particular form of limitation typical of many systems in use today will be discussed in this section, with the object of providing a means of comparing the best results available with these limited systems to the results obtained in the previous sections.

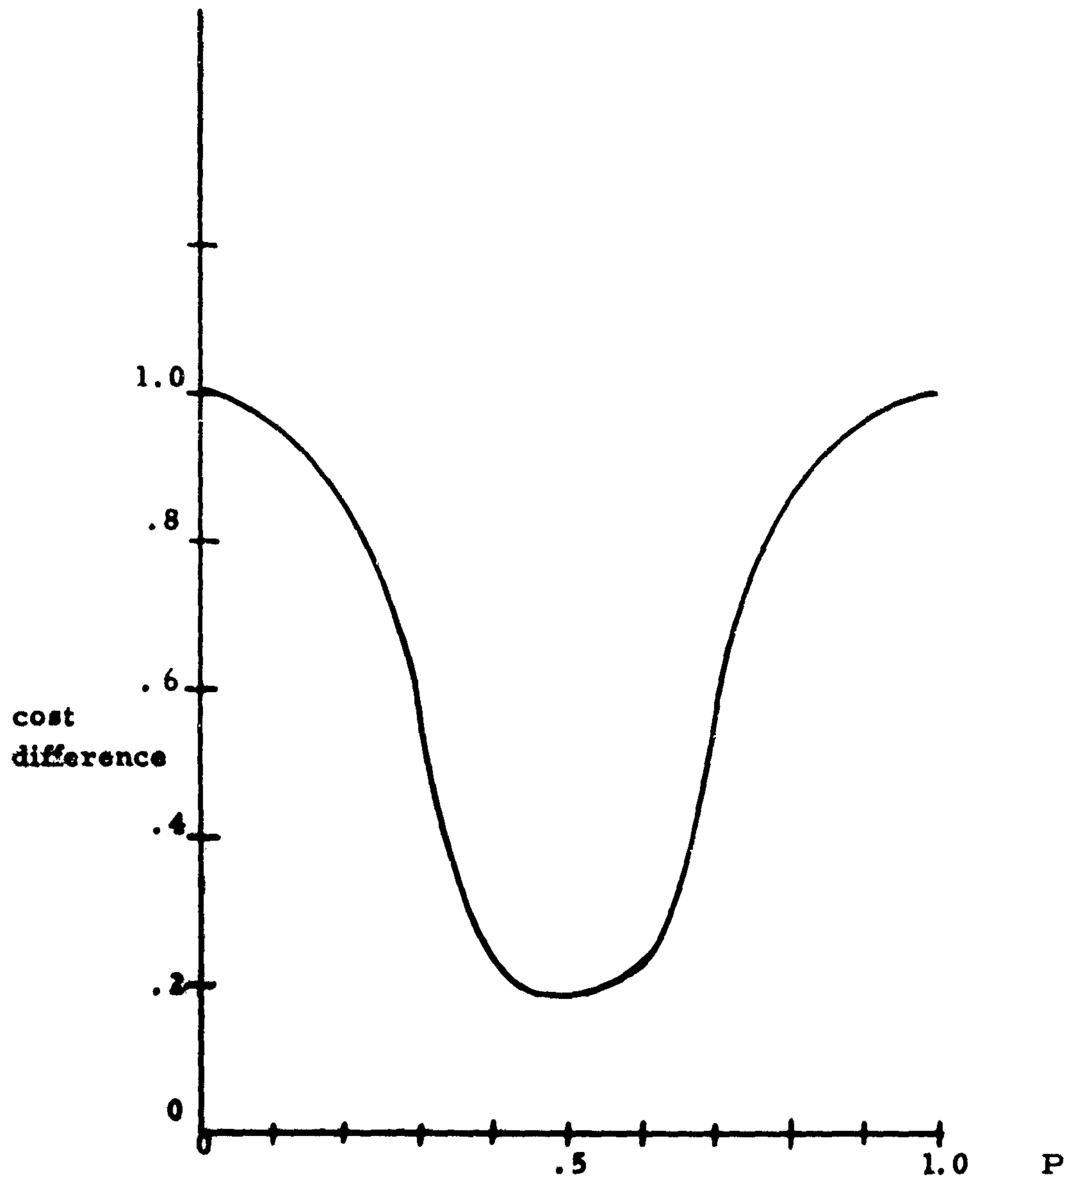

Figure 3.4 Savings in cost over STSD by using optimal sequential strategy

The limitation involved will be the required use of a fixed threshold for every observation, with the output of the observation being simply "over threshold" or "under threshold". An example that illustrates this is the use of an active sonar that every cycle displays either a "blip" or not. Inside the sonar unit there is a discrimination device that can only tell whether or not the received signal statistic is over or under some pre-set threshold, and a blip is displayed or not as a consequence. The observer then must make decisions on the basis of observing a series of binary variables rather than a series of continuous ones.

It is convenient here to let "indication" mean the exceeding of the fixed threshold by a signal. This word must be used with some care, in that an indication alone does not imply any decision concerning the presence or absence of the target. An indication is just a way of reporting the output of a particular (threshold) detection device. Decisions are to be made on the basis of (perhaps) many indications.

The problem facing us now is essentially identical to that stated in section (3.1), with the substitution of a binary random variable y for the continuous x, with probability mass functions $h_0(y)$ and $h_1(y)$ replacing to $p_0(x)$ and $p_1(x)$. If we let y take the values 0 or 1, we have

$$h_0(y) = \begin{cases} 1-f & y=0 \\ f & y=1 \end{cases}$$

$$h_1(y) = \begin{cases} 1-d & y=0 \\ d & y=1 \end{cases}$$

where f and d are the "false indication" and "detection" probabilities defined by

$$f = f(x^*) = \int_{x^*}^{\infty} p_0(x)\,dx$$

$$d = d(x^*) = \int_{x^*}^{\infty} p_1(x)\,dx \qquad (3.12)$$

where x^* is the threshold that is assumed to be fixed throughout the length of the search. (Note that these equations are identical to equations (2.7) for p_f and p_d. We use "f" and "d" for notational ease, and to remind us that they now refer to "indications" in the analysis of this more general decision structure.)

For a given fixed setting of x^*, let us define $F_n(P)$ to be the minimum expected cost obtained by using the optimal truncated fixed threshold strategy. Then by using arguments identical to those leading up to equation (3.10) we may write

$$F_n(P) = \begin{cases} (1-P)\,C_{10} & \delta_n \le P \le 1 \\ P\,C_{01} & 0 \le P \le \gamma_n \\ P\,W_1 + (1-P)\,W_0 + [P\,d + (1-P)f]\,F_{n-1}\left[\dfrac{P\,d}{P\,d+(1-P)f}\right] + & \\ \quad + [P(1-d)+(1-P)(1-f)]\,F_{n-1}\left[\dfrac{P(1-d)}{P(1-d)+(1-P)(1-f)}\right] & \gamma_n \le \end{cases} \qquad (3.13)$$

where we have made use of the fact that

prob. $\{y=1\} = P\,d + (1-P)f$

prob. $\{y=0\} = P(1-d) + (1-P)(1-f)$

$$\text{prob. } \{S_1 | y=1\} = \frac{Pd}{Pd + (1-P)f}$$

$$\text{prob. } \{S_1 | y=0\} = \frac{P(1-d)}{P(1-d) + (1-P)(1-f)} ,$$

It is important to repeat the fact that although there is a fixed detection threshold x^* that leads to indications, there still remain the important points γ_n and δ_n which define the decision regions. These points are referred to as decision thresholds.

In the limit as $n \to \infty$, if we can show that $F_n(P)$ converges, we may call $\lim_{n \to \infty} F_n(P) = F(P)$.

That $F_n(P)$ converges, along with δ_n and γ_n, may be shown exactly as in theorems A, B, C and D of section 3.3.1 with the minor changes in notation necessary to allow for the discrete character of y.

The $F(P)$ and associated strategy that is calculated are still conditioned upon the threshold x^*, in that both d and f are functions of x^*. For a final overall cost minimization, we may select, for every P, that x^* which minimizes $F(P)$. Calling this overall optimum $F_{min}(P)$ we see that

$$F_{min}(P) = \min_{x^*} [F(P)] . \tag{3.14}$$

Quantization of the continuous target information into a binary variable is sure to produce a loss in the information sense, but in the cost context of the present problems we can readily compute the loss in more meaningful cost terms.

Successive iterations of equation (3.13) are again straightforward, and computer calculations are greatly speeded by the absence of an integral. Figure 3.5 shows a typical convergence to $F(P)$ for a particular setting of x^*. Figure 3.6 shows the application of equation (3.14). To avoid confusion in drawing, only some selected curves for $F(P)$ are shown, but in general $F_{min}(P)$ is the lower envelope of all $F(P)$ curves. It is also convenient (and illustrative) to label the $F(P)$ curves by the d corresponding to their x^* through equation (3.12). The values of the parameters have been selected to provide comparison with the $C(P)$ of Figure 3.2, which is shown in dotted lines. The optimal threshold setting (and equivalent d) that results from this example is plotted as a function of P in Figure 3.7. Note that for P outside some range there is no optimal threshold setting, as the decision rule in those cases calls for a terminal decision with no observation taken.

The implementation of the search is identical to that described by section 3.4, where $L(x)$ is replaced by $L(y)$

$$L(y) = \begin{cases} \dfrac{1-d}{1-f} & y = 0 \\[2ex] \dfrac{d}{f} & y = 1 \end{cases} .$$

3.6.1 Numerical Example

On occasion, the shepherd of our previous examples has to use a less versatile meter, one that contains a built-in threshold (that may be set at the start of the search). The signal and noise statistics are still characterized by the p.d.f.s.

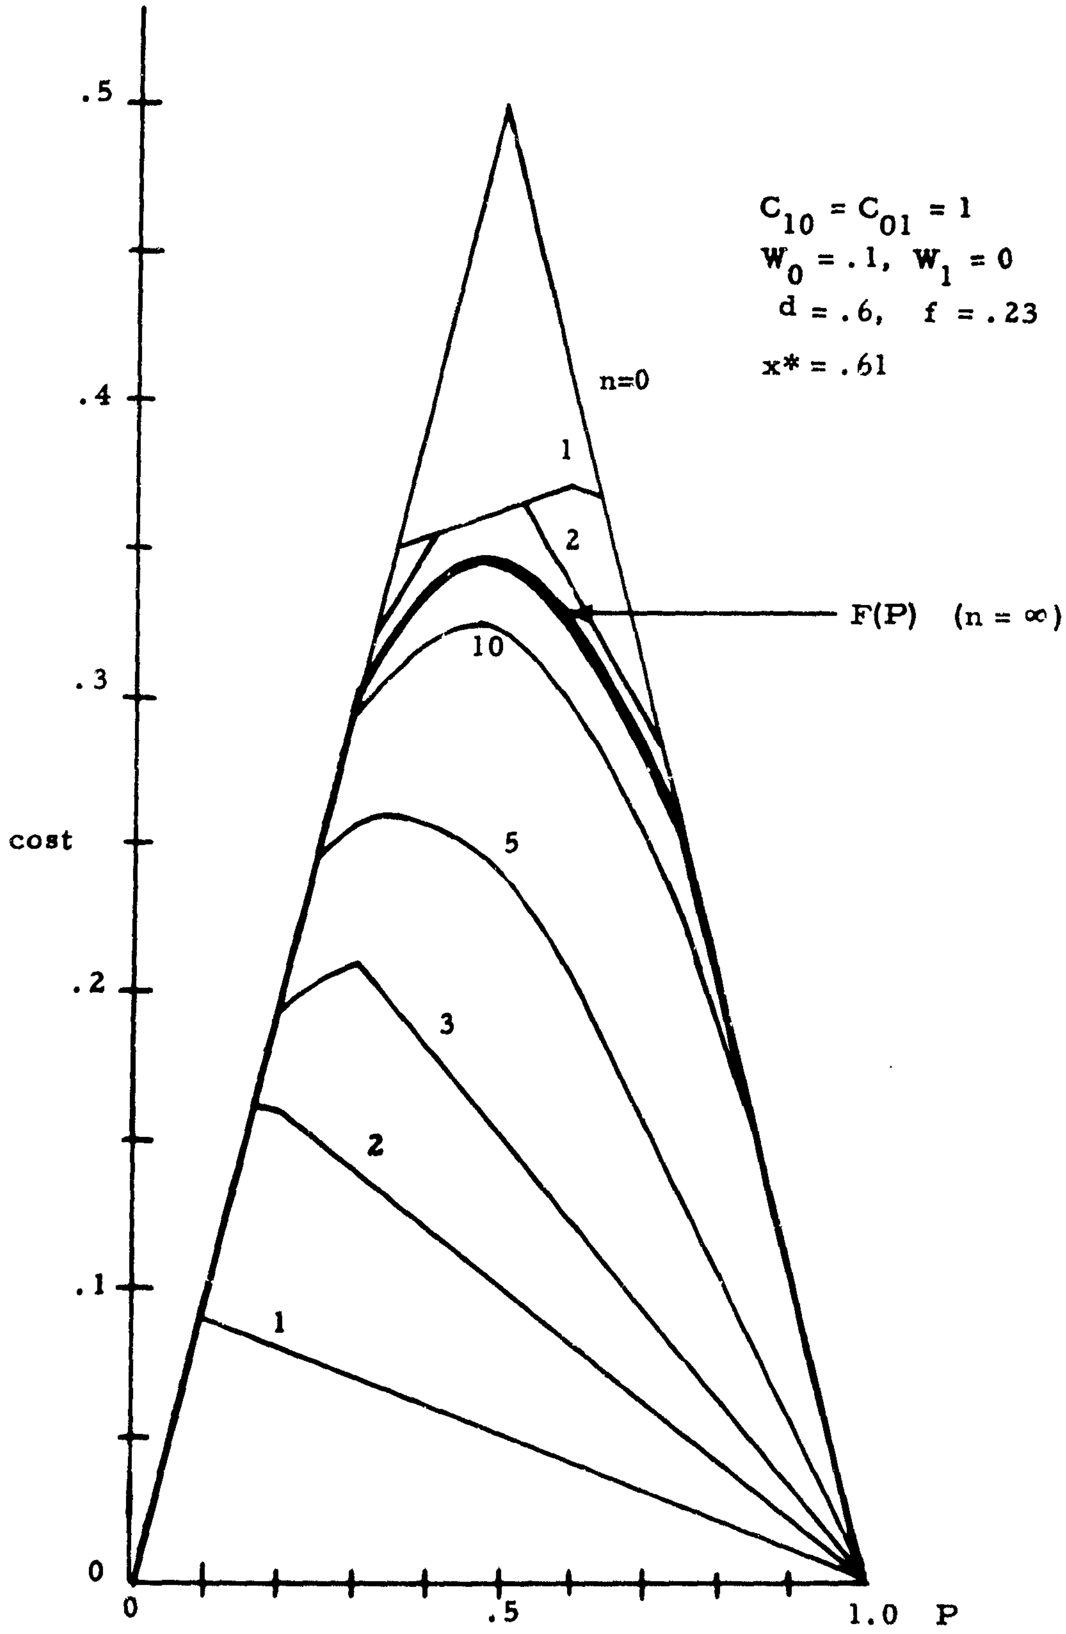

Figure 3.5 Convergence of $F_n(P)$ to $F(P)$

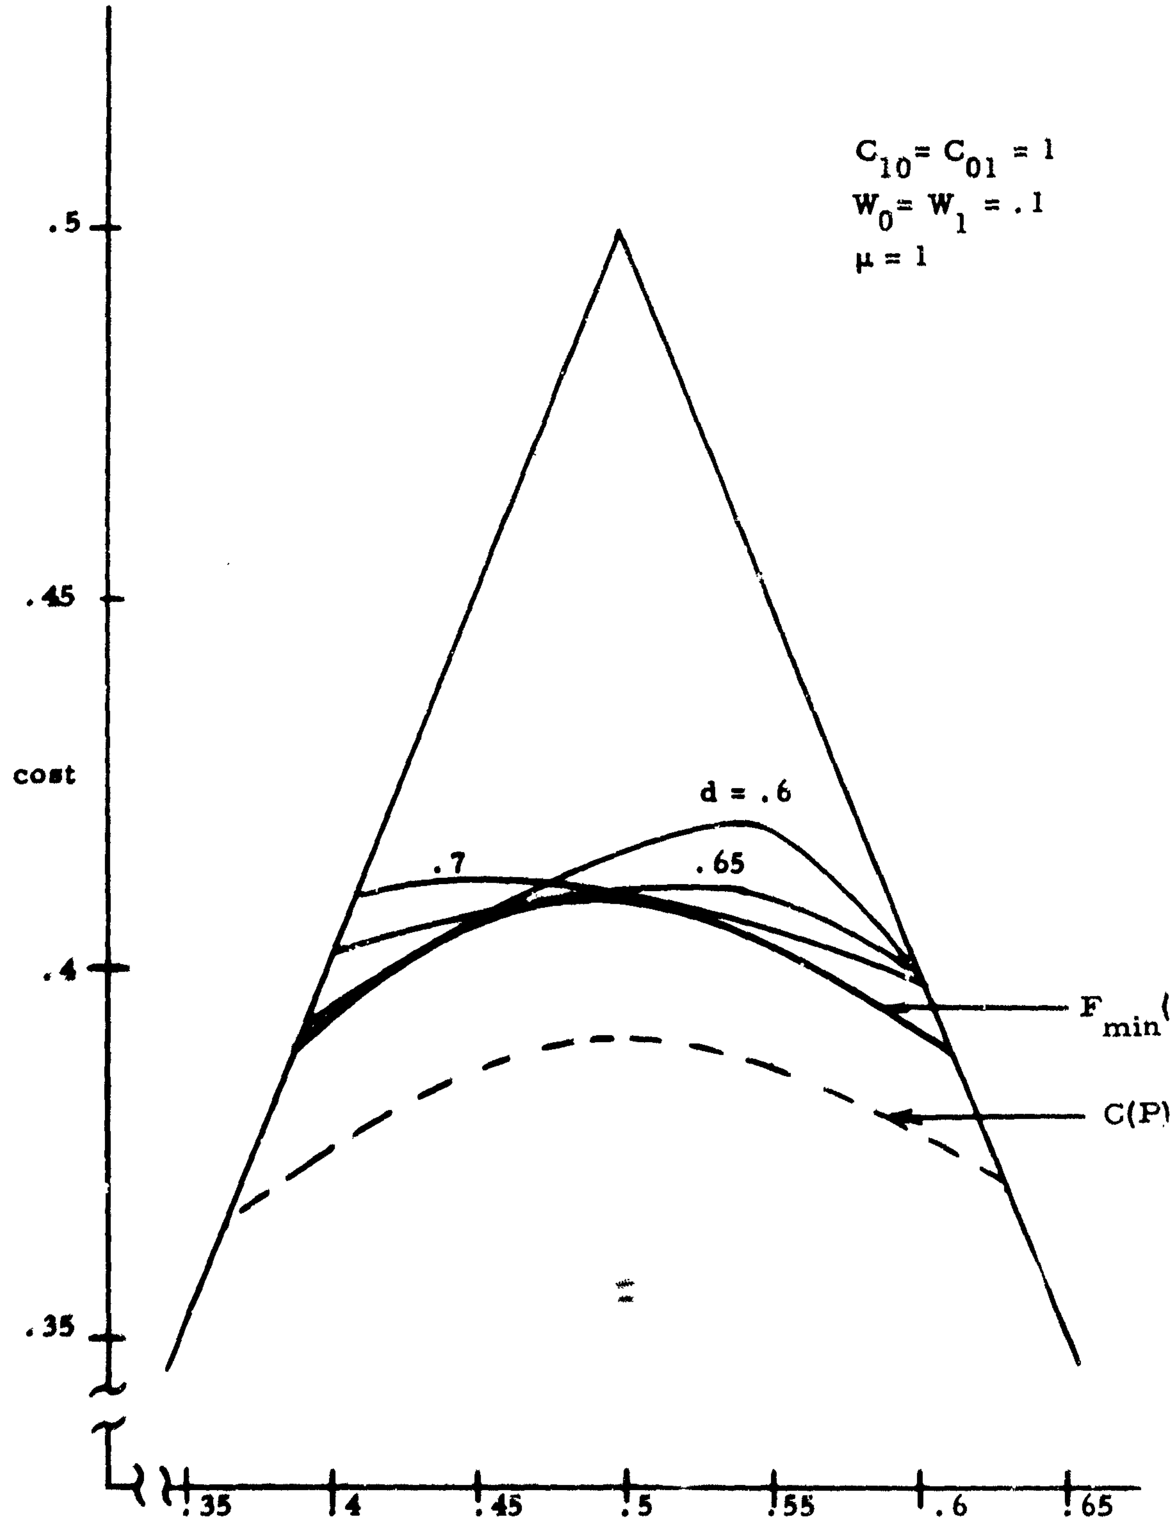

Figure 3.6 Calculation of $F_{min}(P)$

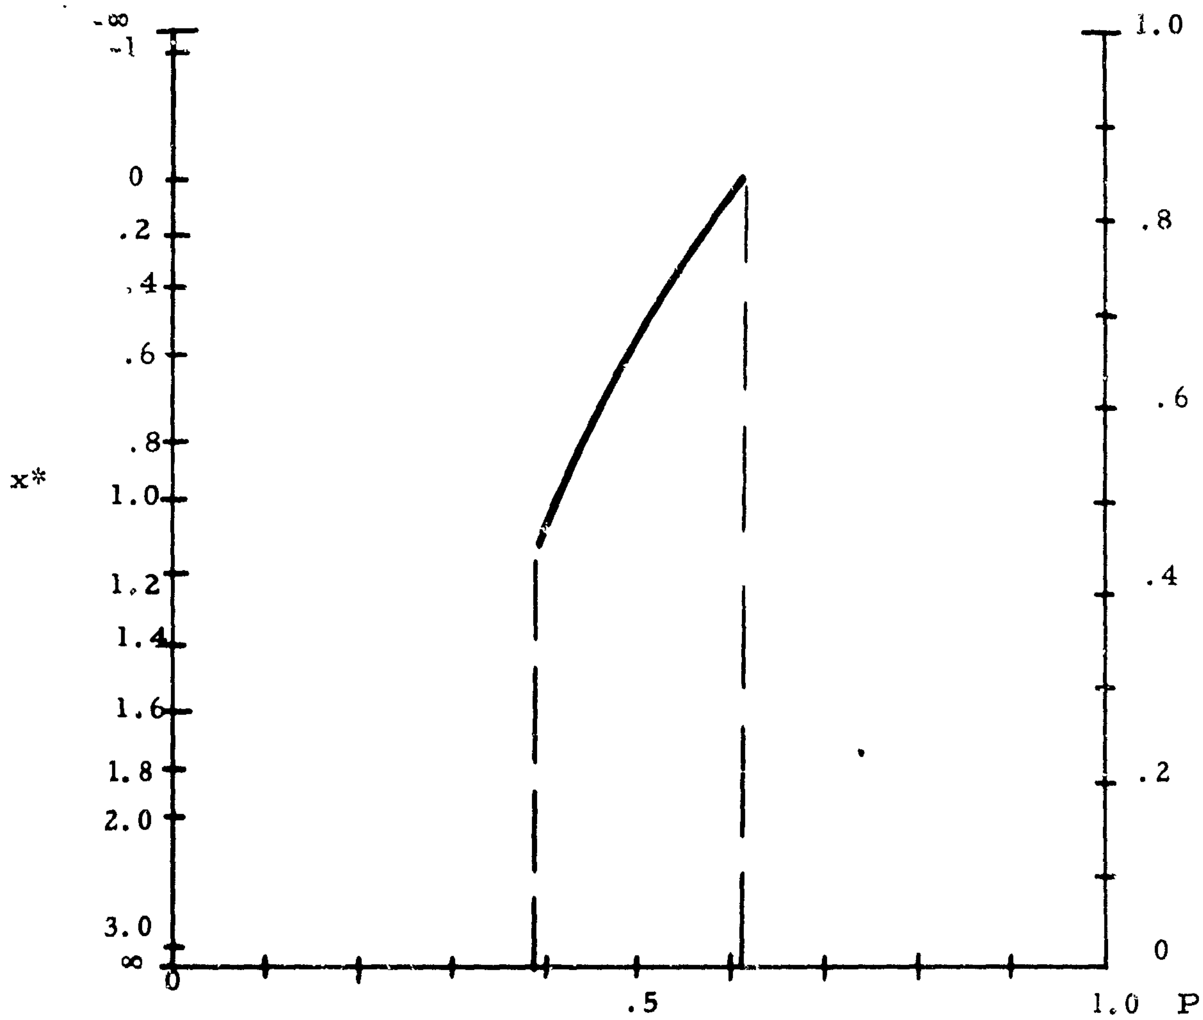

Figure 3.7 Optimal x* as a function of P

$$p_0(x) = f_N(x;\ 0, 1)$$

$$p_1(x) = f_N(x;\ 1, 1)$$

so that, upon setting a threshold x^*, the false indication and detection probabilities will be

$$f = 1 - \mathrm{erf}(x^*)$$

$$d = 1 - \mathrm{erf}(x^* - 1)\ .$$

With the cost values $C_{01} = C_{10} = 1$, $W_0 = W_1 = .1$ of the earlier examples, the shepherd uses Figures 3.6 and 3.7 to obtain the following decision rule:

1. If $P > .61$, don't take any measurements at all, and go directly to the pasture.
2. If $P < .39$, don't take any measurements at all, and go directly to town.
3. If $.39 \leq P \leq .61$, set the threshold as indicated in Figure 3.7, then take a measurement of y.
4. If $y = 0$ $(x < x^*)$ then compare $\frac{1-d}{1-f}$ with the values $\frac{\gamma}{(1-\gamma)} \frac{(1-P)}{P}$ $.64(1-P)/P$, and $\frac{\delta}{(1-\delta)} \frac{(1-P)}{P} = 1.56(1-P)/P$. If $y = 1$ $(x \geq x^*)$ then compare d/f with these values. If the likelihoods fall outside these limits, make the appropriate terminal decision. If within them, take another measurement, and so on.

If, as before, $P = \frac{1}{2}$ we see that at least one measurement must be taken. From Figure 3.7 we find that the threshold is set such that $d = .68$, $f = .3$. If $y = 1$ is observed, since $L(1) = .68/.3 = 2.27$ is greater than 1.56, a D_1 decision is required. If $y = 0$ is observed, since

$L(0) = .32/.7 = .46$ is less than .64, a D_0 decision is required. This strategy (for $P = 1/2$) is idential to the STSD strategy, and so we would expect the search cost at this P to be the same as the STSD cost, which it is. However, at other values of P, the advantage of the sequential threshold strategy becomes apparent.

Illustrations of the difference in costs using the three strategies mentioned so far in this chapter are shown in Figure 3.8. The curves show the percent increase in cost (over the minimum cost attainable by the optimum sequential search) due to using either the non-sequential STSD strategy, or the sequential fixed threshold strategy. As shown by the figures, the importance of using a sequential search is considerable when the experimental cost decreases.

3.7 Adaptive Threshold Sequential Search

The previous section considered the threshold x^* to be fixed throughout the search. This limitation is often attributable to a lack of the time needed to adjust the threshold (if necessary) inbetween possible observation periods. In some systems, however, although the detection device is by nature a threshold indicator, it is possible to vary the threshold inbetween observations. A practical example is the use of a radar PPI scope by an experienced operator. The operator usually keeps the gain of the scope low (to avoid "snow") until a possible blip shows up at some point. The operator can then increase the gain on the next sweep because he will be concentrating on a smaller region of the scope, (i.e. the neighborhood of the possible blip) and so is not as effected by the increased noise.

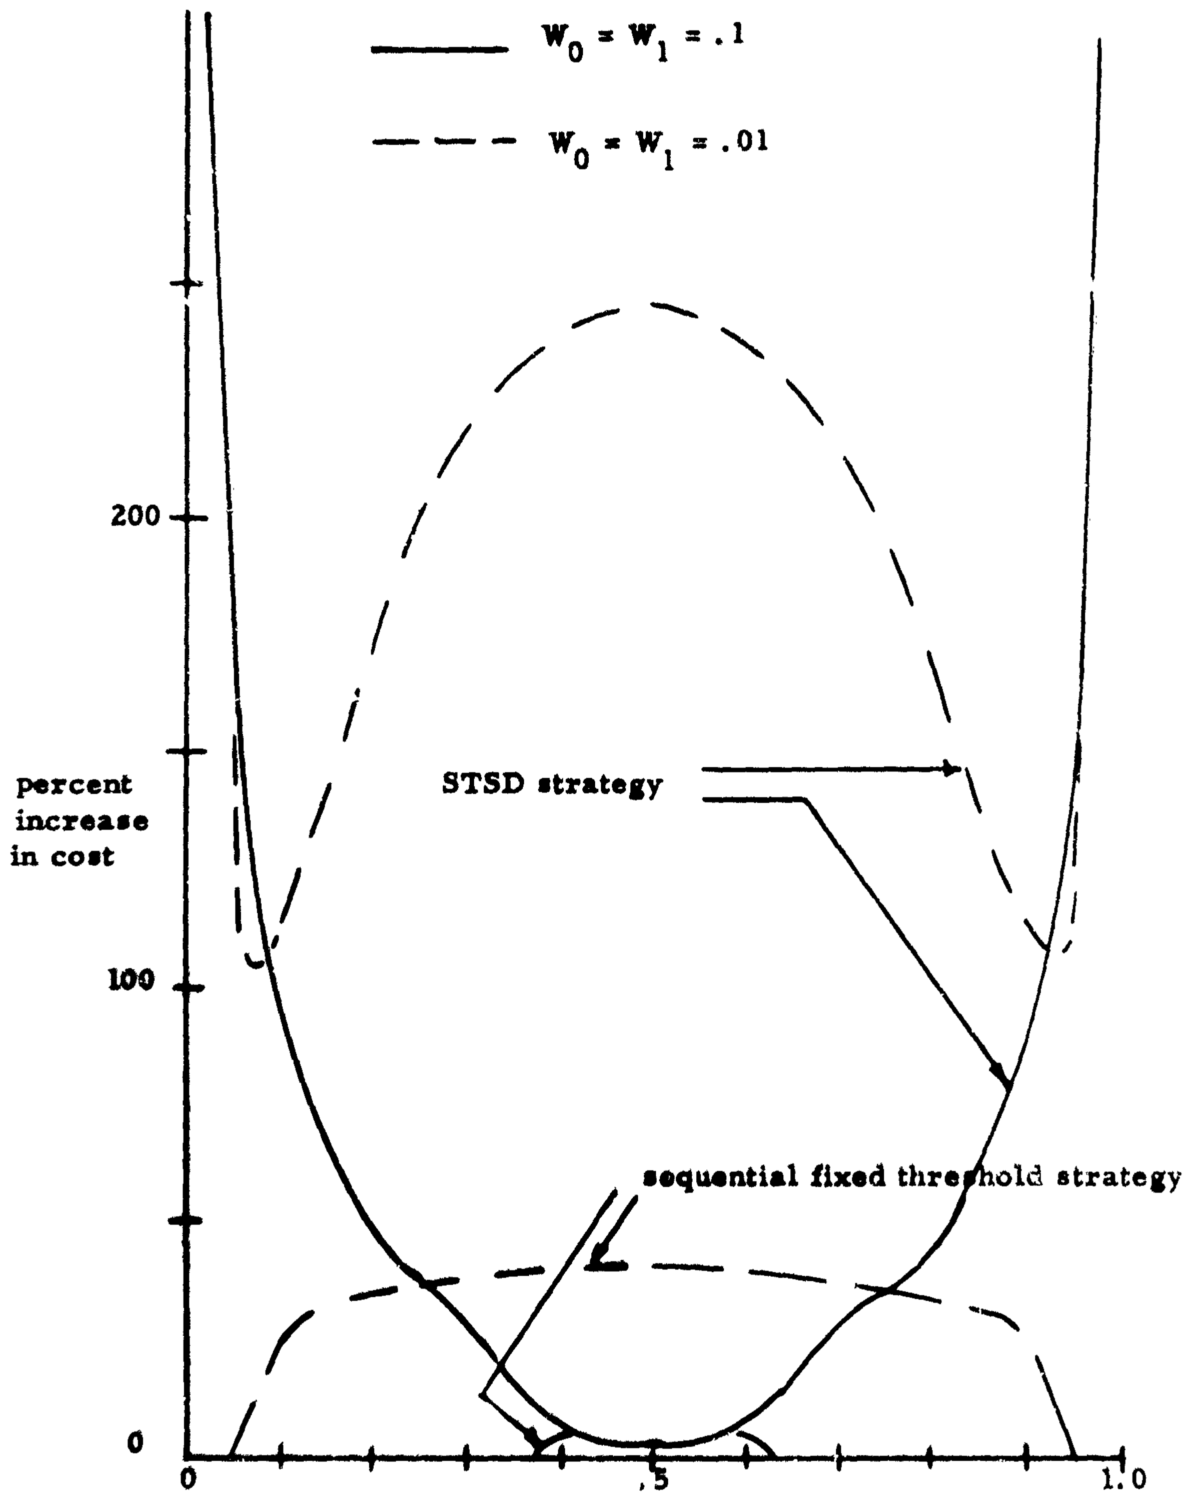

Figure 3.8 Percent increase in cost using non-optimal strategy

This process of controlling the threshold is often referred to as the "alerting effect" and shows up in our model of sequential search in the following way. Let the decision structure be such that the searcher is required to take two observations, and that he must make a terminal decision after the second observation, both limited to threshold observations. If he determines the first threshold x_1^* on the basis of STSD, he has (with $C_{00} = C_{11} = 0$)

$$L(x_1^*) = \frac{1-P}{P}\left(\frac{C_{10}}{C_{01}}\right).$$

The resultant first observation can either be $y = 1$ or $y = 0$ (indication, or no indication). Now suppose $y = 1$ is observed. Then the posterior "target present" probability P' is

$$P' = \text{prob.}\{S_1 | y=1\} = \frac{P p_d}{P p_d + (1-P) p_f} .$$

He would now want to set the second threshold so that

$$L(x_2^*) = \frac{1-P'}{P'}\left(\frac{C_{10}}{C_{01}}\right) = \frac{p_f}{p_d}\left(\frac{1-P}{P}\right)\frac{C_{10}}{C_{01}} = \frac{p_f}{p_d}\; L(x_1^*) .$$

Since in practical systems $p_f < p_d$, we see that $L(x_2^*) < L(x_1^*)$. With the usual monatonic likelihood function, this implies that $x_2^* < x_1^*$, so the second threshold has been lowered. Because this lowering of the threshold between first and second observations yields a corresponding tendancy towards deciding the target is present, the expression "alerting" is descriptive. The "alert" is due to the reception of $y = 1\,(x_1 > x_1^*)$. If $y = 0\,(x_1 < x_1^*)$ were received, there would be an equivalent "dulling" effect by an increase in the second threshold.

This sort of adaptive threshold adjustment may be carried over into the framework of the previous problems, and calculations can be made to compare the resulting minimum cost with the solutions of the fixed threshold and continuous variable searches of the previous sections. In order to show this, let us define $K_n(P)$ to be the minimum cost attained by using an optimal truncated adaptive threshold strategy. Then we may write

$$K_n(P) = \min \begin{cases} T(P) \\ \min\limits_{x^*} L_n(P, x^*) \end{cases}$$

where

$$L_n(P, x^*) = PW_1 + (1-P)W_0 + [Pd + (1-P)f]\, K_{n-1}\left(\frac{Pd}{Pd+(1-P)f}\right)$$
$$+ [P(1-d) + (1-P)(1-f)]\, K_{n-1}\left(\frac{P(1-d)}{P(1-d)+(1-P)(1-f)}\right)$$

The dependence of $L_n(P, x^*)$ on x^* is a consequence of $d = d(x^*)$, $f = f(x^*)$ as defined in equations (3.12).

Again, calculation of $K_n(P)$ is straightforward, but rather lengthy due to having to explore the x^* variation at each stage of the iteration, rather than just at the limit as in the previous section. The approach of $K_n(P)$ to the limit function $K(P)$ as $n \to \infty$ is assured by simple modifications to the proofs given in section 3.3.1. The resulting minimum cost function and strategy is again similar to those derived before.

Although calculations of $K(P)$ have not been carried out, it is possible to prove the intuitive inequality

$$C(P) \leq K(P) \leq F_{min}(P) \quad \text{for all } 0 \leq P \leq 1$$

the proof is not instructive, and is omitted here.

CHAPTER IV

TARGET ARRIVAL AT A RANDOM TIME

4.1 Towards a More Realistic Model of Certain Searches

All the previous work in this paper has been carried out with the condition that the states of nature describing the system were fixed throug out the search. This was a basic part of the model first described in Chapter II, and it is only on this basis that any of the hypothesis tests, sequential or otherwise, could be applied. In the analysis of sequential search strategies, this condition of stationarity of the target can be seen to be a very limiting one, in that the duration of the search is a random variable. Thus "throughout the search" means all that time for which the probability of the search duration is non-zero. This time range is often infinite. The use of a hypothesis test must therefore be carefully considered, and should be used <u>only</u> when the only alternatives to the states of nature are either S_0 and S_1, and when only one of these states will definitely hold for the entire search.

Many practical problems which have been analysed from the point of view of hypothesis testing are much more reasonably approached by a new model which will be treated in this chapter. This model allows the target to arrive in the cell at some time t (a random variable), after the start of the search and then remain there for the rest of the search. The decision rules should then be based upon the probability that the target has arrived yet, rather than whether or not it is present at all.

This sort of model is certainly applicable in many military detection models, and the term "raid recognition" has often been used to describe the general problem statement. When searching for a submarine in a particular area, for example, we rarely have the luxury to assume that it will either be there or not for the rest of the search (i.e. time of interest). In particular, when the decision D_0 = "Target not present" is made, the search equipment would not be turned off. In fact, in problems of this type the decisions become limited to two decisions: D_1 or W. Let us continue the submarine search example to develop typical actions and cost factors.

Suppose decision $D_1(\tau)$ is the order to send ASW aircraft at time τ to the region representing the cell covered by the detection device. If the submarine has not yet arrived, then the aircraft must return to base, some resultant cost of false decision has been incurred, and the search continues. If the submarine has arrived at some time t (before the D_1 decision was made) then the longer the delay between t and τ the more difficult an eventual interception, and thus there is an increase in what we might call the interception cost when the $D_1(\tau)$ decision is made. The search is then considered to be terminated. One obvious objective is to make decisions so that the expected cost of the search is minimized. This chapter will consider such a model, and the optimal search that evolves from it.

We shall also spend some effort in the development of a simple way of comparing some specific non-optimal rules that are being practiced or proposed, by use of a concept similar to the ROC described in Chapter II.

4.2 General Problem Statement

1. The target arrives in the cell of interest at time t ($t=0, 1, \ldots$) with known probability $p(t)$.

2. If the cell is observed at time τ, the result is a random variable x which has p.d.f. $p_0(x)$ if $t > \tau$ and $p_1(x)$ if $t \leq \tau$.

3. At every time τ ($\tau=0, 1, \ldots$) the searcher makes one of the following decisions

   $D(\tau)$: Decide the target has arrived

   $W(\tau)$: Wait for another observation.

4. The decision $D(\tau)$ may or may not be a terminal decision: If $D(\tau)$ is picked and $t > \tau$, then a false decision cost F is incurred, the knowledge that $t > \tau$ is gained, and the search continues.

   If $D(\tau)$ is picked and $t \leq \tau$, then the search is terminated with a cost of $\phi(t, \tau)$.

5. The objective is to minimize the expect cost of the search. The strategy achieving this minimum cost is called the "optimal" strategy.

Although we shall not solve the general problem as stated above, it is helpful to keep it in mind when solving it subject to reasonable assumptions. Again, as in the previous work, we shall restrict $p_0(x)$ and $p_1(x)$ such that $L(x) = p_1(x)/p_0(x)$ is monotonic non-decreasing in x, to ease the notation.

4.3 Linear Terminal Cost, and Geometrical Arrival Time

Two assumptions will now be made in order to obtain a solution. These also have the advantage of reasonably representing some real search situations.

First, we shall assume that the terminal cost is proportiona to the time "late", i.e. $\phi(t,\tau) = (\tau - t)W$. This form of the function is not necessary for a solution, but offers a minimum of algebraic difficulties that might otherwise cloud the development.

Second, and more restrictive, we shall assume that the targ arrival time distribution is

$$p(t) = \lambda(1-\lambda)^t \qquad t = 0, 1, \ldots \tag{4.1}$$

This geometric distribution has the advantage in describing the arrival as being conceptually "random" by the fact that it provides a constant probability of arrival per unit time (λ), given that it hasn't yet arrive This property also provides a simple representation for a state variab by summing up the total information about the state of nature (i.e. wh or not the target has arrived). This is outlined as follows.

If we let $P(\tau)$ be the probability that the target has arrived at or before time τ, then by equation (4.1)

$$P(\tau) = \sum_{t=0}^{\tau} p(t) = 1 - (1-\lambda)^{\tau+1} \tag{4.2}$$

Furthermore, we can show that

$$P(\tau+1) = (1-\lambda)P(\tau) + \lambda$$

so that (without any other information) to describe $P(\tau+1)$, all that is needed is $P(\tau)$.

In addition, suppose that a value of $x(\tau)$ is observed at time τ. Then

$$P(\tau+1 \mid x(\tau)) = \frac{\text{prob. } (x(\tau), \tau+1)}{\text{prob. } x(\tau)}$$

To calculate prob. $(x(\tau), \tau+1)$, we note that this could happen in two ways: the target could have arrived at time τ or before, in which case $p_1(x)$ is the p.d.f. of $x(\tau)$, or the target could have arrived at time $\tau+1$, so that $p_0(x)$ is the p.d.f. Thus

$$P(\tau+1 | x(\tau)) = \frac{[P(\tau)p_1(x) + [1-P(\tau)]\lambda p_0(x)]}{P(\tau)p_1(x) + (1-P(\tau))p_0(x)} \tag{4.3}$$

and we see that given some observation x , the posterior probability of the target arriving at or before time $\tau+1$ is still dependent only upon $P(\tau)$, and not τ explicitly. It is this basic Markov property that allows us to proceed now in a way similar to the approach in Chapter III (where, we recall, the Markov property of the successive likelihood ratios led to the dynamic program approach).

The assumption leading to equation (4.1) also enables us to characterize the start of the search, since $P(0) = \lambda$. What is more important, we note that according to statement (4) of the problem definition, when the false decision $\{D(\tau) | t>\tau\}$ is made, knowledge that $t > \tau$ is automatically gained. This fact and equation (4.1) lead to $P(\tau+1 | t>\tau) = \lambda$.

We are now prepared to write a functional equation for the minimum search cost. Let us define $V(P(\tau))$ as the minimum search cost obtained using the optimal search strategy at time τ where $P(\tau)$ is the present value of the probability that the target has arrived previous to or at time τ . There are two decision choices. One is $D(\tau)$: decide target has arrived, with resultant probability being wrong of $1-P(\tau)$, and subsequent cost of F plus what the continued search will cost from then on. The other is $W(\tau)$: wait for more information x ,

in which case the "late" cost is incurred only if the target has arrived (probability $P(\tau)$), and the search continues with the proper posteriori probability given by equation (4. 3). The minimum cost is then given by

$$V(P) = \min \begin{cases} (1-P)(F+V(\lambda)) & : D \\ PW+\int_{-\infty}^{\infty}[P p_1(x)+(1-P)p_0(x)]\, V\left[\dfrac{P p_1(x)+(1-P)p_0(x)\lambda}{p_1(x)+(1-P)p_0(x)}\right] & : W \end{cases} \tag{4.4}$$

The τ has now been purposely left out as an argument of P and the decisions, since this equation holds for all τ and only $P(\tau) = P$ is needed to express the right-hand side. In what follows τ will be left out except when necessary to avoid ambiguity.

One result is immediately apparent. By letting $P=0$ we have

$$V(0) = \min\,[F + V(\lambda), V(\lambda)] = V(\lambda) \quad \text{for} \quad F > 0 .$$

This tells us that the cost of search, if we know the target has not arrived, is the same as if we waited one time unit and started again. This is because there can be no "late" cost W if the target has not yet arrived.

In order to develop a feeling for the solution to equation (4. 4), and to obtain an upper bound upon $V(P)$, the next section shall consider the degenerate case getting no information from the observations.

4.4 Optimal Search With No Information

Suppose that $p_0(x) = p_1(x)$. Then, as can be shown by equation (4.3), an observation of x does not affect the posterior evaluation of P . In this case the observation x is irrelevent, and the searcher gains no information. The searcher may still develop an optimal strategy, which now consists simply of either waiting one time unit, or deciding the target has arrived. Equation (4.4) becomes

$$V(P) = \min \begin{cases} (1-P)(F + V(\lambda)) & : D \\ \\ WP + V[P + (1-P)\lambda] & : W \end{cases} \tag{4.5}$$

As noted before, $V(0) = V(\lambda)$, and we can also easily see that $V(1) = 0$.

From the form of equation (4.5) it is postulated that the structure of the strategy will be

$$\text{if } P \geq \gamma : D$$
$$P \leq \gamma : W$$

where γ is a decision point to be determined as part of the solution.

That this is indeed the form of the strategy, and that it is not degenerate (that is, $0 < \gamma < 1$), may be shown by the following proof by contradiction. (The discussion that follows can also be shown to be valid for the more general equation (4.4). Since it is the form of the proof that is of interest, it is carried out in this less complicated case.)

Let us define

$$D(P) = (1-P)(F + V(\lambda))$$

$$G(P) = WP + V[P + (1-P)\lambda] \quad .$$

A sketch of these functions is shown in Figure 4.1. $D(P)$ is a straight line with $D(0) = F + V(\lambda) = F + V(0)$, $D(1) = 0$. $G(P)$ has an unknown functional dependence on P through $V[P + (1-P)\lambda]$, but it is continuous by the continuity of $V(P)$. The boundary values are known and are $G(0) = V(\lambda) = V(0)$, $G(1) = W + V(1) = W$. Since $F > 0$ and $W > 0$, then $G(1) > D(1)$ and $G(0) < D(0)$ so that $G(P)$ and $D(P)$ must intersect at at least one point.

Suppose $G(P)$ were such that $G(P)$ and $D(P)$ intersect at more than one point (in Figure 4.1 this is illustrated by the dotted line), say γ' , γ'' and γ''' . Let us select a point P' such that

$$\gamma' < P' < \gamma''$$

$$\gamma'' < P' + (1-P')\lambda < \gamma''' \tag{4.6}$$

where $[\gamma', \gamma'']$ is a D region and $[\gamma'', \gamma''']$ is a W region. Then by equation (4.5)

$$V(P') = (1-P')(F + V(\lambda)) < WP' + V[P' + (1-P')\lambda] \tag{4.7}$$

and

$$V(P' + (1-P')\lambda) = W(P' + (1-P')\lambda) + V[P' + (1-P')\lambda + (1-P')(1-\lambda)\lambda] <$$
$$< (1-P')(1-\lambda)[F + V(\lambda)] \quad . \tag{4.8}$$

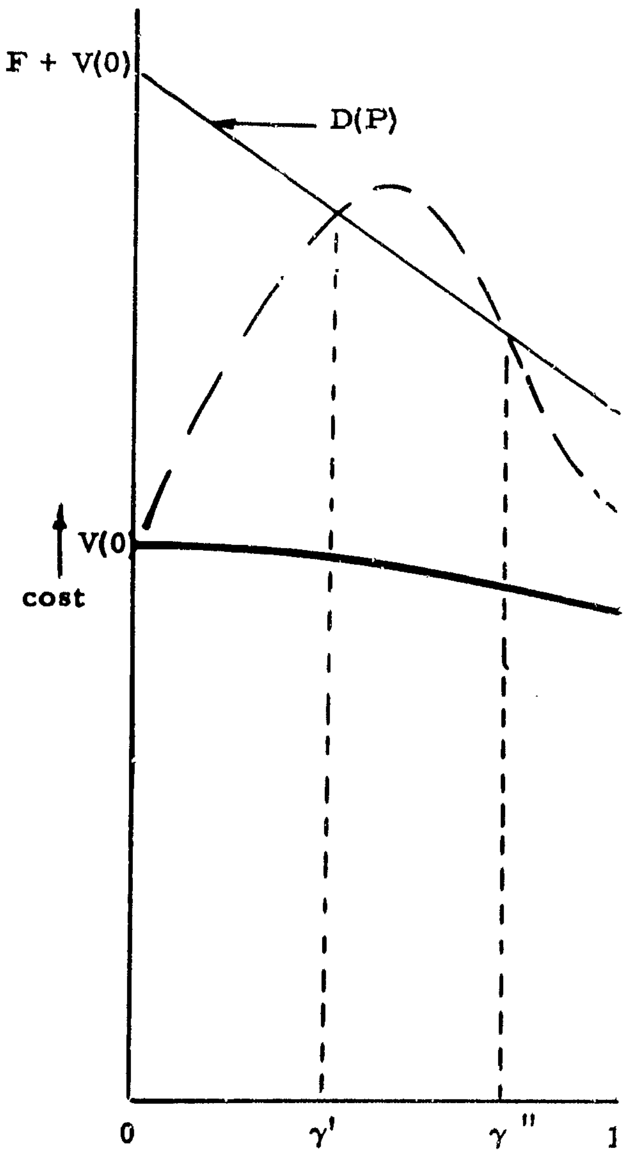

Fi

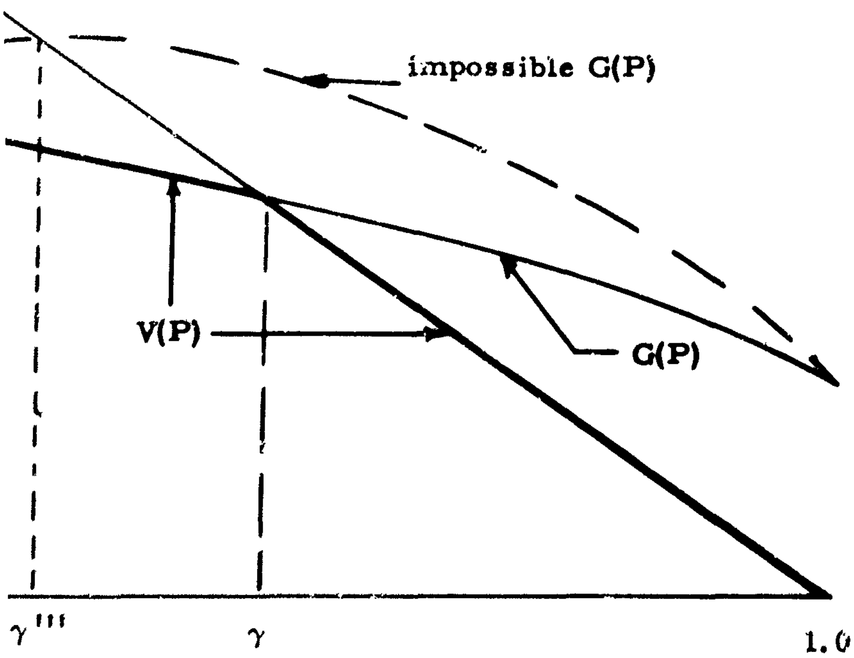

The construction of V(P)

Combining equations (4.7) and (4.8)

$$(1-P')(F + V(\lambda)) < WP' + (1-P')(1-\lambda)[F + V(\lambda)]$$

which reduces to

$$P' > \frac{\lambda (F + V(\lambda))}{W} \quad . \tag{4.9}$$

If we select another point P'' so that

$$\begin{aligned} &\gamma'' < P'' < \gamma''' \\ &\gamma''' < P'' + (1-P'')\lambda < 1 \end{aligned} \tag{4.10}$$

we can show in a similar manner that

$$P'' < \frac{\lambda (F + V(\lambda))}{W}$$

which, with equation (4.9) implies that $P'' < P'$. But, since equations (4.6) and (4.10) require $P' < \gamma'' < P''$, the contradiction is proven.

We have just shown then that $G(P)$ and $D(P)$ intersect at one point, and defining this point as $P = \gamma$, we re-write equation (4.5)

$$V(P) = \begin{cases} (1-P)(F + V(\lambda)) & P \geq \gamma \\ WP + V[P + (1-P)\lambda] & P \leq \gamma \end{cases} \tag{4.11}$$

As a first step in the solution of this equation, let $P = \gamma$. Then since $\gamma + (1-\gamma)\lambda > \gamma$ we have $V(\gamma) = (1-\gamma)(F+V(\lambda)) = W\gamma + V(\gamma + (1-\gamma)\lambda) = W\gamma + (1-\gamma)(1-\lambda)(F+V(\lambda))$ from which we get

$$F + V(\lambda) = \frac{\gamma}{1-\gamma} \frac{W}{\lambda} \tag{4.12}$$

so that equation (4.11) becomes

$$V(P) = \begin{cases} (1-P)\frac{\gamma}{1-\gamma}\frac{W}{\lambda} & P \geq \gamma \\ \\ WP + V[P + (1-P)\lambda] & P \leq \gamma \end{cases} \tag{4.13}$$

The next step is to find γ in terms of W, F and λ. Once this is obtained, the optimal strategy is defined. (Determination of the functional form of $V(P)$ for $P \leq \gamma$ will then rely upon iterations of (4.13) in a manner to be described later.)

To determine γ, let us assume that γ has been obtained and is such that

$$\begin{aligned}
&\lambda < \gamma \\
&1-(1-\lambda)^2 < \gamma \\
&\quad\vdots \\
&1-(1-\lambda)^{n-1} < \gamma \\
&1-(1-\lambda)^n \geq \gamma
\end{aligned} \tag{4.14}$$

where n is the smallest integer such that equation (4.14) holds. By n-1 successive applications of equation (4.13) we get

$$\begin{aligned}
V(\lambda) &= W\lambda + V[1-(1-\lambda)^2] \\
&= W\lambda + W(1-(1-\lambda)^2) + V[1-(1-\lambda)^3] \\
&\ \ \vdots \\
&= W[(n-1) - \frac{1-\lambda}{\lambda}(1-(1-\lambda)^{n-1})] + V[1-(1-\lambda)^n] ,
\end{aligned}$$

all (n-1) steps being the result of W-decisions. Finally, since the n^{th} must be a D-decision:

$$V(\lambda) = W[(n-1) - \frac{1-\lambda}{\lambda}(1-(1-\lambda)^{n-1})] + (1-\lambda)^n \frac{\gamma}{1-\gamma}\frac{W}{\lambda} .$$

Using the value of $V(\lambda)$ from equation (4.12) we may solve for γ in terms of n

$$\gamma = 1 - \frac{1-(1-\lambda)^n}{\lambda\,(\frac{F}{W}+n)} \tag{4.15}$$

By use of equation (4.14) we find that n is the smallest integer such that

$$(1-\lambda)^n \leq \frac{1}{1+\lambda\,(\frac{F}{W}+n)} \tag{4.16}$$

Once n is found, γ is then obtained from equation (4.15).

We have just proved that the form of the strategy consists of waiting for a fixed amount of time (number of time units) $n-1$, then choosing a D-decision. In the event that the search has just started, (or that a D-decision has just been made but the target has not yet arrived, so that the search must be resumed with $P = 0$), this fixed amount of time is given by equation (4.16), from which γ can be determined.

If the search starts out so that $P \neq \lambda$, then by successive applications of equation (4.13) we can show in a calculation similar to that above that the time until a D-decision, $n(P)$, is given by the smallest $n(P)$ that satisfies

$$(1-P)(1-\lambda)^{n(P)-1} \leq 1 - \gamma\ . \tag{4.17}$$

In order to compute $V(P)$ we again simply apply equation (4.13), $[n(P)-1]$ times with decision W, and the $n(P)^{th}$ time with a D-decision. This gives the following form of the minimum cost, where $n = n(P)$

$$V(P) = W[(n-1) - \frac{1-P}{\lambda}(1 - (1-\lambda)^{n-1})] + (1-P)(1-\lambda)^{n-1} \frac{\gamma}{1-\gamma} \frac{W}{\lambda} \tag{4.18}$$

When the mean arrival time gets very large (so that $\lambda \ll 1$), an interesting approximation holds. Since n gets large, we may consider equation (4.16) to be an equality, and as $\lambda \to 0$ we find

$$1 + \lambda\left(\frac{F}{W} + n\right) = (1-\lambda)^{-n} \doteq 1 + n\lambda + \frac{n(n+1)}{2}\lambda^2 + \dots$$

so that

$$n(n+1) \cong \frac{2F}{W\lambda}$$

or, since n is very large, $n \cong \sqrt{\frac{2F}{W\lambda}}$. With this approximation equation (4.15) becomes

$$\gamma \cong \frac{\frac{F\lambda}{W} + \sqrt{\frac{2F\lambda}{W}}}{1 + \frac{F\lambda}{W} + \sqrt{\frac{2F\lambda}{W}}}$$

and so $V(0) = V(\lambda) = \sqrt{\frac{2FW}{\lambda}}$ from equation (4.12). Thus, if the search always starts with $P = 0$ or $P = \lambda$, this expression gives the minimum expected search cost.

4.1.1 Numerical Example of the Non-Informative Search

In this section we shall treat a simple example of the search analysed in the previous section. The solution is in itself interesting and is also useful in order to compare the result with an informative search to be treated later.

It is known that an unfriendly trawler will arrive in a certain region of the ocean in order to cut some trans-oceanic cables. The defense of this region is carried out by a unit that can dispatch high speed aircraft to the region which are able to identify and deter the trawler. If, when the aircraft are sent out, the trawler is not present, they note this fact, return to base and a flight cost of \$10,000 is incurred. However, for every half-hour (unit time) period after it arrives that the trawler is unchallenged in the region, it does \$ 1,000 worth of damage. It is assumed that the time that the trawler will arrive at the region is a geometrically distributed random variable with a mean of 5 hours (10 unit time periods), so that $\lambda = .1$. In units of \$10,000, $F = 1$ and $W = .1$.

Using equation (4.16)

$$(.9)^n \leq (2 + .1n)^{-1}$$

we find that the smallest n that satisfies this expression is $n = 11$. Thus the optimal strategy, given that the trawler is not present at the start of the search, is to wait (n-1) time units (5 hours) before sending out the aircraft, and repeat this procedure, until the trawler is found.

A calculation of γ is made with equation (4.15)

$$\gamma = 1 - \frac{1 - (.9)^{11}}{2.1} = .673$$

and the expected cost at the beginning of the search is

$$V(0) = V(\lambda) = \frac{\gamma}{1-\gamma} \frac{W}{\lambda} - F = \frac{.673}{.327} - 1 = 1.06 \ .$$

If P is the probability that the trawler is in the region at the start of the search, then $n(P)$ and $V(P)$ may be calculated by mean of equations (4.17) and (4.18) with the value of $\gamma = .673$. A plot of $V(P)$ is shown in Figure 4.2. Note that $V(P)$ consists of the lower envelope of 11 straight lines. These represent the cost of search if the strategy is to wait $(n-1)$ units of time before sending the aircraft, where $n = 1, 2, \ldots, 11$. The n for the minimum cost at any P is thus $n(P)$. For example, for $P \geq \gamma$, $n = 1$ and equation (4.18) gives

$$V(P) = \frac{\gamma}{1-\gamma} \frac{W}{\lambda} (1-P) = 2.06(1-P) \quad .$$

4.5 Previous Work Relating to the Non-Informative Search

Before considering the informative search model it should be noted that the non-informative problem has been considered in the literature, but not within the framework of search. Barlow et. al (2) discuss this problem as an example of a "checking" procedure, and consider its applications particularly to problems of checking equipment that is subject to random failure. Thus the D-decision is the decision to check the part to see if it has failed yet, with some appropriate nuisance cost if it hasn't, while the W-decision is the decision to wait one more time unit. They do not attempt to treat the extension of the problem to the possibility of observing some noisy signal associated with the failing part (a temperature reading, for example). The problem is defined in terms of continuous times, so that n is replaced by t , and equation (4.14) and equations derived from it become equalities.

What is an interesting comment on their work, however, is the fact that although they assume a general form of the arrival distribution, they never prove the form of the optimal solution. What is done, in effect, is to assume that the optimal checking procedure will be to wait some time t_1 , then check, if no failure is seen, wait some time t_2 , check, etc. They go on to show that for the exponential failure time density function (the equivalent of our geometric mass function) these t_i are equal. The assumption of the (now proven correct) fixed checking time form

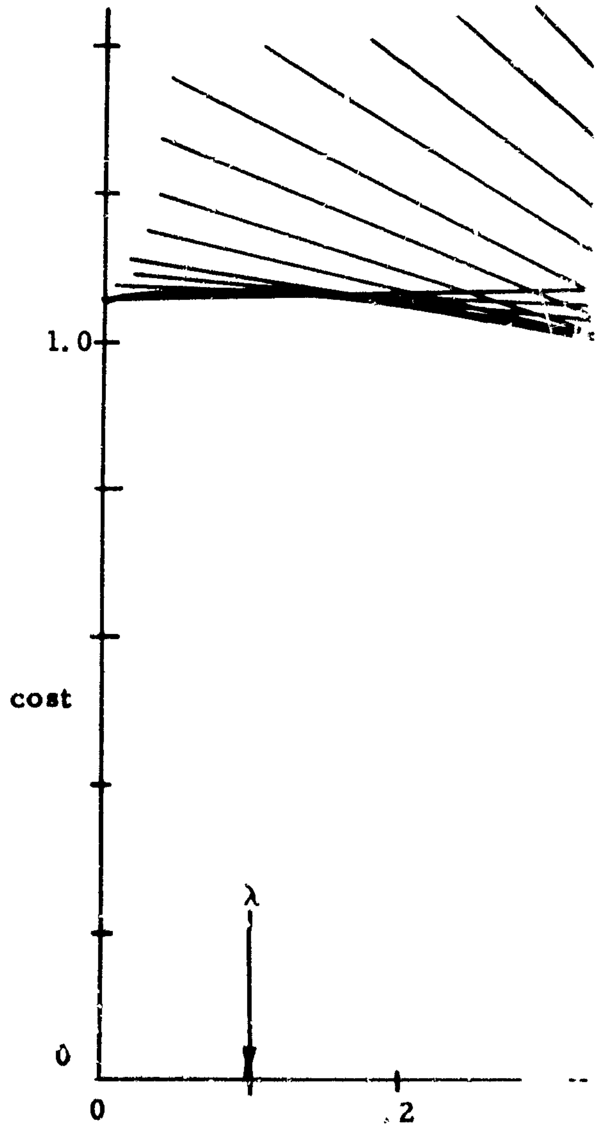

Figure

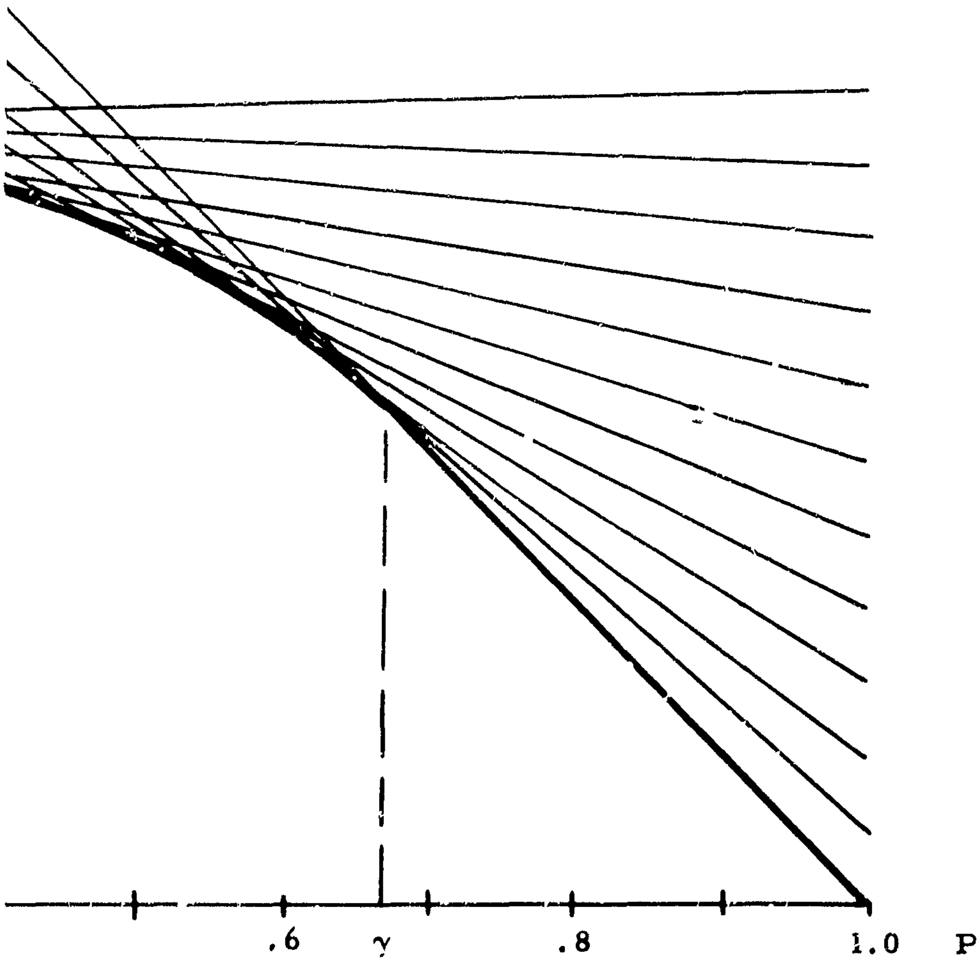

P) with no information

of solution enables them to immediately write a cost expression in terms of the checking time t , which can be minimized by differentiation.

The dynamic programming approach adopted in section 4.2 thus has the advantage of providing the form of the optimal strategy, as well as its parameters. In addition, as will be seen below, we are now in a position to simply extend the structure to allow for observations of a random variable relating to the state of nature (whether or not it concerns a target arriving, or a piece of failing equipment).

4.6 Optimal Search With Information

We are now ready to attack the problem involving observations of a random variable x with every W decision. First we note that the proof in section 4.4 concerning the form of the non-informative search may be carried over conceptually to the more general case represented by equation (4.4), so we state without formal proof that equation (4.4) may be written

$$V(P) = \min \begin{cases} D(P) \\ H(P) \end{cases} = \begin{cases} D(P) & P \geq \gamma \\ H(P) & P \leq \gamma \end{cases} \tag{4.19}$$

where

$$D(P) = (1-P)(F + V(\lambda))$$

$$H(P) = PW + \int_{-\infty}^{\infty} [P\,p_1(x) + (1-P)\,p_0(x)]\; V\left(\frac{P\,p_1(x) + (1-P)\,p_0(x)\,\lambda}{P\,p_1(x) + (1-P)\,p_0(x)}\right) dx$$

We are now faced with a functional equation similar to that treated in Chapter III. Unfortunately, however, there is no efficient way to intuitively truncate the search, as was done in that case. In spite

of this lack of intuitive truncation, it is of course still possible to solve equation (4.19) by such an iterative procedure. In fact, a standard technique for solving such functional equations, and many transcendental equations, is simply the method of successive approximations used in the previous section. If convergence properties can be shown, then any such method is valid, despite the non-physical character of intermediate solutions.

For this reason, we shall re-write equation (4.19) with V as a function of an iteration index n (that has no particular connection to any physical index of the search). Doing this yields

$$V_n(P) = \min \begin{cases} D_n(P) = (1-P)(F + V_{n-1}(\lambda)) \\ H_n(P) = PW + \int_{-\infty}^{\infty} [P p_1(x) + (1-P) p_0(x)] \; V_{n-1}\left(\frac{P p_1(x) + (1-P) p_0(x)}{P p_1(x) + (1-P) p_0(x)} \right) \end{cases} \tag{4.20}$$

All that is needed now is the selection of the boundary condition $V_0(P)$, and assurance that successive iterations will converge the process to $V(P)$ as $n \to \infty$. If we let $V_0(P) = 0$ for all P, then

$$V_1(P) = \min [(1-P)F \, , \; PW] \geq 0 = V_0(P) \quad .$$

With the fact that we have found some $V_n(P) \geq V_{n-1}(P)$, a proof very similar to that leading to theorem B of section 3.3.1 allows us to show that in fact all $V_n(P) \geq V_{n-1}(P)$, so that the process will approach $V(P)$ from below. To complete the convergence, we need to show that $V(P)$ is bounded from above. This can be shown by noting that

$$V(\lambda) \leq (1-\lambda)(F + V(\lambda))$$

so that $V(\lambda) \leq F \frac{1-\lambda}{\lambda}$ and therefore

$$V(P) \leq (1-P)(F + V(\lambda)) \leq \frac{1-P}{\lambda} F \leq \frac{F}{\lambda} \tag{4.21}$$

Again, as in Chapter III, it is easy to show that γ_n, the solution of

$$D_n(\gamma_n) = H_n(\gamma_n)$$

converges to some limit γ as $n \to \infty$.

An example of this iteration process is shown in Figure 4.3, where we have selected the values of the parameters to compare to the results of section 4.4.1. The $p_0(x)$ and $p_1(x)$ are normal distributions with unit variance and mean of 0 and μ respectively, where μ is the signal to noise ratio.

These calculations were calculated on an IBM 7090. As can be seen from the figure, the convergence is much slower than that for the problems of Chapter III. However, the computations are still essentially additions, with an appropriate approximation for the integral.

The general solution of the search for a randomly arriving target, with observations of an appropriate signal, has thus been obtained, with the minimum cost attainable and optimal strategy comprising the solution. As pointed out in a previous section, this model applicable to a randomly arriving target may also be applied to such problems as a randomly failing piece of equipment.

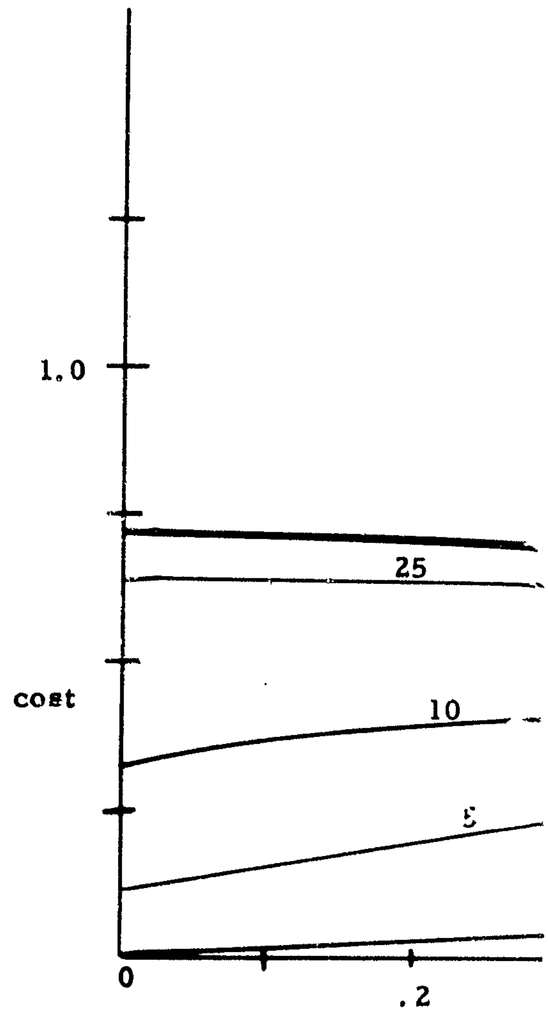

Figure 4. 3

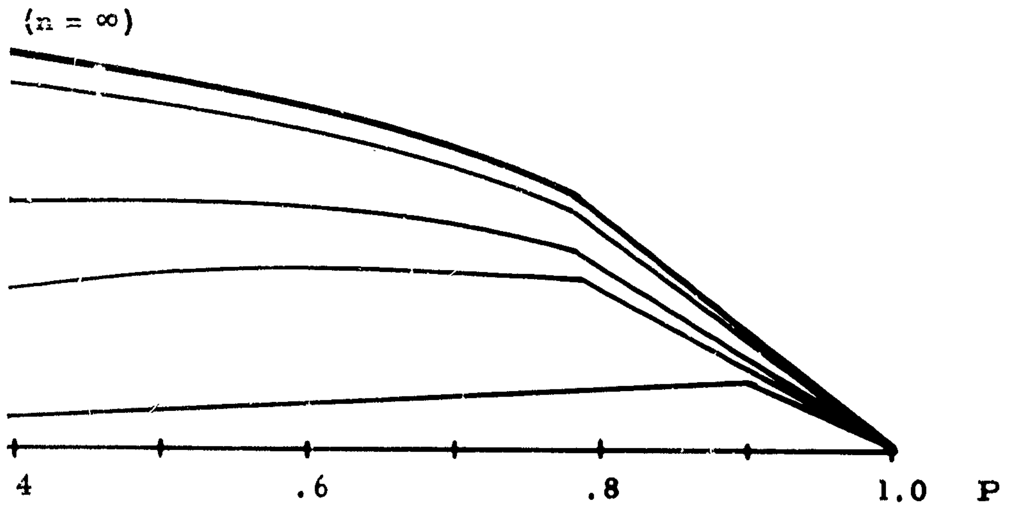

ce of $V_n(P)$ to $V(P)$ with information

4.6.1 Numerical Example of Informative Sequential Search

The defense unit of section 4.4.1 has decided to install a sonar apparatus to aid in the search against the trawler. The results of integrating sonar signals for the half-hour unit time is a random variable x that has p.d.f. $p_0(x)$ if the trawler has not yet arrived, or $p_1(x)$ if the trawler has arrived, where $p_0(x) = f_N(x;\ 0, 1)$, $p_1(x) = f_N(x;\ 1, 1)$. As indicated in section 3.5.1, this is the case when detecting a known signal . additive Gaussian noise, with a signal-to-noise ratio $\mu = 1$.

From Figure 4.3 we note that $\gamma = .78$. This is higher than the value of .673, and indicates that the availability of information will let the searcher be less quick to respond. We also note that $V(0) = V(\lambda) \cong .58$, which is a saving of close to 50% compared to the non-informative search. The strategy that gives these results follows.

Suppose $P = \lambda = .1$ to start the search. Since $.1 \leq \gamma = .78$, an observation is required at the first time interval. Suppose a value x_1 is the result of this observation.

The posteriori probability $P(x_1)$ that the target has arrived is now given by equation (4.3)

$$P(x_1) = \frac{(.1)\,p_1(x_1) + (.1)(.9)\,p_0(x_1)}{(.1)\,p_1(x_1) + (.9)\,p_0(x_1)} = \frac{(.1)\exp(x_1 - \frac{1}{2}) + .09}{.1\ \exp(x_1 - \frac{1}{2}) + .9}$$

Comparing this with γ, we see that if

$$\exp(x_1 - \tfrac{1}{2}) \geq 27.8$$

or

$$x_1 \geq .5 + \ln 27.8 \approx 5.9$$

then the aircraft should be sent out. If not, another observation x_2 should be made, the a posteriori probability $P(x_1, x_2)$ should be calculated and compared to γ, etc.

4.6.2 A Comment on The Solution

As is shown by the example in Figure 4.3, although convergence of $V_n(P)$ to $V(P)$ is guaranteed, the speed with which the process converges is rather slow. In fact, as λ gets very small, causing the cost of search to increase, the convergence is even slower. This unfortunate practical difficulty is at present unresolved. One possible approach is suggested here.

We decided in section 4.6 to start the iteration with $V_0(P) =$ which consequently assures convergence from below. It is equally possible to start the iteration at some appropriate large value, which will again assure convergence, but then from above. One such value would be the right hand side of the condition given in equation (4.21), i.e. $V_0(P) = \frac{F}{\lambda}$. However, a lower starting point is available by noting that the minimum cost of the informative search is less than or equal to the minimum cost of the non-informative search, for all values of P. This lower starting value of $V_0(P)$ could considerably decrease the number of iterations needed to provide a given degree of accuracy.

Other techniques for establishing a reasonable first guess of $V(P)$, and letting this equal $V_0(P)$, would be a valuable.

aid in the computation. In general, however, convergence proofs might be difficult for arbitrary starting functions $V_0(P)$.

4.7 Implementation of the Strategy and Comments On the Geometric Arrival Assumption

In Chapter III we showed that the use of the optimal search strategy resulted in essentially a Wald sprt, where the decision boundaries were determined by cost considerations rather than by error probabilities. A similar analysis of the implementation of the strategy developed in section 4.6 is of interest, in that it points out a basic limitation to the treatment of the problem.

From the form of the decision structure presented in equation (4.18), we see that P , the probability that the target has arrived up to some time, is constantly compared to some decision threshold γ . When a D-decision occurs, P automatically returns to 0 if the target has not yet arrived. With a series of W-decisions, however, a series of observations $x_1, x_2, \ldots$ has been made, and the posterior probability of the target having arrived can be derived.

Specifically, let us consider the search to start with a W-decision at $\tau = 1$, and that n successive observations of $x_1, x_2, \ldots, x_n$ are made. We shall also consider a completely general arrival time distribution $f(t)$, $t=1, 2, \ldots$ (we define $f(0) = 0$) . Using the statement of conditional probability, where $\underline{x} = (x_1, x_2, \ldots, x_n)$, we define

$$P_n = \text{prob.}\{t \leq n | \underline{x}\} = \frac{\text{prob.}\{\underline{x} | t \leq n\}\ \text{prob.}\{t \leq n\}}{\text{prob.}\{\underline{x}\}} .$$

The unconditional probability of receiving some vector $\underline{x}$ is

$$\text{prob.}\{\underline{x}\} = f(1)\prod_{i=1}^{n} p_1(x_i) + f(2)\, p_0(x_1)\prod_{i=2}^{n} p_1(x_i) + f(3)\, p_0(x_1)\, p_0(x_2)\prod_{i=3}^{n} p_1$$

$$+ \ldots f(n)\left(\prod_{i=1}^{n-1} p_0(x_i)\right) p_1(x_n) + \sum_{j=n+}^{\infty} f(j)\left(\prod_{i=1}^{n} p_0(x_i)\right)$$

$$= f(1)\prod_{i=1}^{n} p_1(x_i) + \sum_{j=2}^{n} f(j)\left(\prod_{i=1}^{j-1} p_0(x_i)\right)\left(\prod_{k=j}^{n} p_1(x_k)\right) + \sum_{j=n+1}^{\infty} f(j)\left(\prod_{i=1}^{n} p_0(\cdot\right.$$

(4.22)

Similarly, we may calculate

$$\text{prob.}\{\underline{x}|t\le n\} \cdot \text{prob.}\{t\le n\} = f(1)\prod_{i=1}^{n} p_1(x_i) + \sum_{j=2}^{n} f(j)\left(\prod_{i=1}^{j-1} p_0(x_i)\right)\left(\prod_{k=j}^{n} p_1(x_k\right.$$

(4.23)

The condition for a D-decision, given that $\underline{x}$ has been observed, is that $P_n \ge \gamma$. Using the above expressions this condition becomes

$$f(1)\prod_{i=1}^{n} p_1(x_i) + \sum_{j=2}^{n} f(j)\left(\prod_{i=1}^{j-1} p_0(x_i)\right)\left(\prod_{k=j}^{n} p_1(x_k)\right) \ge \frac{\gamma}{1-\gamma}\sum_{j=n+1}^{\infty} f(j)\left(\prod_{i=1}^{n} p_0(x_i)\right).$$

Dividing by $\prod_{i=1}^{n} p_0(x_i)$ gives the condition

$$\sum_{j=1}^{n} f(j)\left(\prod_{i=j}^{n} \frac{f_1(x_i)}{f_0(x_i)}\right) = \sum_{j=1}^{n} f(j)\left(\prod_{i=j}^{n} L(x_i)\right) \ge \frac{\gamma}{1-\gamma}\sum_{j=n+1}^{\infty} f(j) \qquad (4.24)$$

where $L(x_i)$ is the likelihood ratio for each observed x_i .

We note now an interesting relation between this search and the one developed in Chapter III, and described in section 3 4. If we assume that $f(t)$ is such that

$$f(1) = P$$

$$f(j) = 0 \qquad j < \infty$$

$$f(\infty) = 1 - P$$

(that is, the target either arrives at $t=1$ with probability P , or it never arrives so that $f(\infty) = 1-P$) , then the condition (4.24) becomes

$$\prod_{i=1}^{n} L(x_i) \geq \frac{\gamma}{1-\gamma} \frac{1-P}{P}$$

which is precisely the condition needed in the sprt to reach a D_1 decision. We see then that we have not only a sprt, but one with no lower decision point, which implies that $W_0 = 0$ (see the proof of theorem D and the definition of Γ). We have shown then that a certain class of the problems described by equation (3.2), namely those for which $W_0 = 0$, are imbedded in the general solution of equation (4.4).

Returning to the problem at hand, in the general case represented by equation (4.24), we can no longer describe the test that develops as a simple random-walk. In fact the equation represents a very complicated process. Not only is the $\prod_{i=j}^{n} L(x_i)$ term weighted by the $f(j)$, this weighting is successively compared to a term which gets smaller as n increases. Since the simple

random-walk with constant absorbing barriers has not been fully solved (see Appendix A), there is no reason to believe that this non-Markovian (because P_n is more than just a function of P_{n-1}) process with non-constant barrier would be any easier. Thus a strictly Wald-type approach depending as it does upon the statistics of such a process, would not seem too profitable.

If we let $f(t) = \lambda (1-\lambda)^{t-1}$ $(t = 1, 2, \ldots)$, however, an interesting result is shown. Equation (4.24) becomes

$$\sum_{j=1}^{n} \lambda (1-\lambda)^{j-1} \prod_{i=j}^{n} L(x_i) \geq \frac{\gamma}{1-\gamma} \sum_{j=n+1}^{\infty} \lambda (1-\lambda)^{j-1} = \frac{\gamma}{1-\gamma} (1-\lambda)^n$$

and by defining $\Lambda(x_i) = L(x_i)/(1-\lambda)$ we get

$$\sum_{j=1}^{n} \prod_{i=j}^{n} \Lambda(x_i) \geq \frac{\gamma}{(1-\gamma)\lambda}$$

which has the advantage of being a test that compares a variable

$$Z_n = \sum_{j=1}^{n} \prod_{i=j}^{n} \Lambda(x_i)$$

to a constant decision threshold.

In addition, the sequence Z_n describes a Markov process in that Z_{n+1} only depends upon Z_n (as well as $\Lambda(x_{n+1})$, of course). To show this we note that

$$Z_{n+1} = Z_n \Lambda(x_{n+1}) + \Lambda(x_{n+1})$$

which can be verified by direct substitution into the definition of Z_n.

The geometric distribution of arrival times thus imparts a Markov character to the decision process. And indeed, it is just this character that has allowed us to approach the problem from the dynamic programming point of view. By allowing the argument of V to be P (and similarly the argument of C in Chapter III), we have been assuming that P is completely descriptive of the searcher's state of knowledge about the system, and that the history of events that lead to P are unimportant. Conversely, since for the general arrival time distribution it cannot be shown that equation (4.24) represents a Markov process, we cannot write a general equation similar to equation (4.4) as P alone is not sufficient to represent a "state" of the process.

4.8 Fixed Threshold Sequential Search

As a prelude to the next sections, we shall consider here the optimal search with the use of threshold observations defined in section 3.6. Again we shall consider the comparison of x to some fixed threshold x^* which remains constant throughout the search. Defining, as before

$$f = f(x^*) = \int_{x^*}^{\infty} p_0(x)\,dx$$

$$d = d(x^*) = \int_{x^*}^{\infty} p_1(x)\,dx$$

we may write an equation similar to equation (4.4), with $R(P)$ here representing the minimum expected cost obtained by using the optimal fixed threshold sequential strategy, and the observations consist of $y = 1$ $(x \geq x^*)$ and $y = 0$ $(x \leq x^*)$.

$$R(P) = \min \begin{cases} (1-P)(F + R(\lambda)) \\ \\ PW + [Pd + (1-P)f]\, R\left(\dfrac{Pd + (1-P)f\lambda}{Pd + (1-P)f}\right) + \\ \qquad + [P(1-d) + (1-P)(1-f)]\, R\left(\dfrac{P(1-d) + (1-P)(1-f)\lambda}{Pd + (1-P)(1-f)}\right) . \end{cases} \tag{4.25}$$

As in the previous sections we shall compute $R(P) = \lim_{n\to\infty} R_n(P)$ using the equation

$$R_n(P) = \min \begin{cases} A_n(P) \\ \\ B_n(P) \end{cases} = \begin{cases} A_n(P) & P \geq \gamma \\ \\ B_n(P) & P \leq \gamma \end{cases} \tag{4.26}$$

where

$$A_n(P) = (1-P)(F + R_{n-1}(\lambda))$$

$$B_n(P) = PW + [Pd + (1-P)f]\, R_{n-1}\left(\frac{Pd + (1-P)f\lambda}{Pd + (1-P)f}\right)$$
$$+ [P(1-d) + (1-P)(1-f)]\, R_{n-1}\left(\frac{Pd + (1-P)(1-f)\lambda}{Pd + (1-P)(1-f)}\right)$$

and

$$R_0(P) \equiv 0 .$$

Again, we note that $R(P)$ is a function of x^* through d and f, and so the final optimisation takes place for each P (and, in particular, for $P = \lambda$) such that

$$R_{min}(P) = \min_{x^*} [R(P)] .$$

An example of such a calculation is shown in Figures 4.4, 4.5, and 4.6 with values of W, F, λ and μ similar to those used in previous examples for the no observation case, and for the continuous x case. Figure 4.4 shows the convergence of $R_n(P)$ to $\mathcal{R}(P)$ for a particular x value. Figure 4.5 is a plot of $\mathcal{R}(\lambda)$ as a function of x^* (and thus a function of d), and Figure 4.6 is a plot of $R_{min}(P)$ as a function of P.

4.8.1 Numerical Example of Fixed Threshold Sequential Search

We return to the search for the trawler of the previous example, and assume that the sonar device has a built-in threshold which must be set permanently before the search starts. Given the same cost parameters as before, we see from Figure 4.5 that if the search starts with $P = \lambda$ the threshold x^* must be set to equal $x^* = .88$ so that $d = .55$, $f = .19$. Using this value of the threshold we see from Figure 4.6 that $\gamma = .75$ and $V(\lambda) = .67$, which is a higher cost than the search with observation of continuous x, but a lower cost than obtained in the non-informative search. These other two costs are also illustrated by the other lines in Figure 4.6.

The fixed threshold strategy that results is as follows. Suppose that the search starts with $P = \lambda$. Since $\lambda = .1 < \gamma = .75$, an observation y_1 should be taken. The two possible output results, are $y_1 = 0$ $(x \leq x^*)$ and $y_1 = 1$ $(x \geq x^*)$, which produce the a posteriori probabilities,

$$P = \begin{cases} \dfrac{(.1)(1-d) + (.1)(.9)(1-f)}{(.1)(1-d) + (.9)(1-f)} = .15 & y_1 = 0 \\[2em] \dfrac{(.1)d + (.1)(.9)f}{(.1)d + (.9)f} = .32 & y_1 = 1 \end{cases}$$

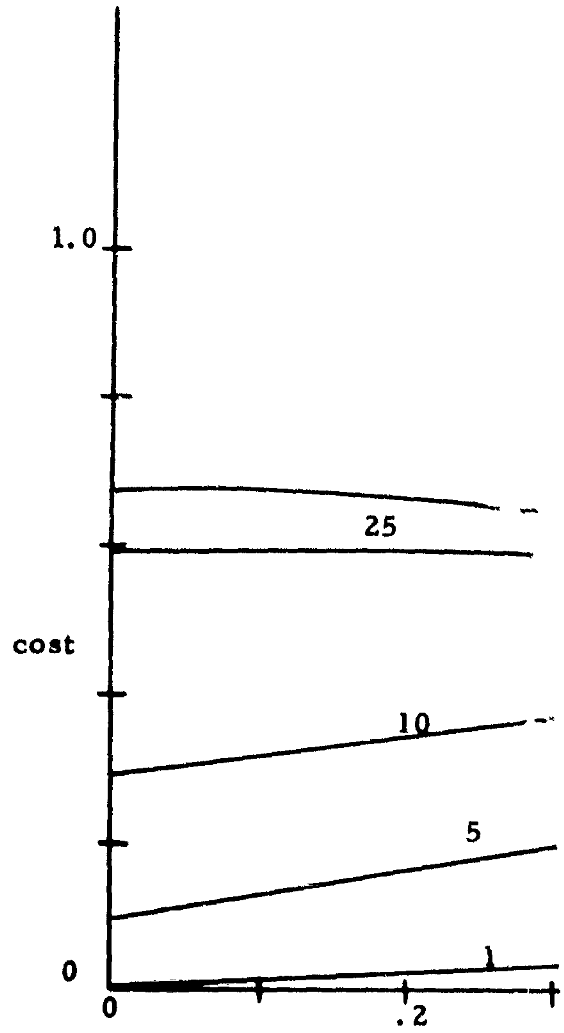

Figur

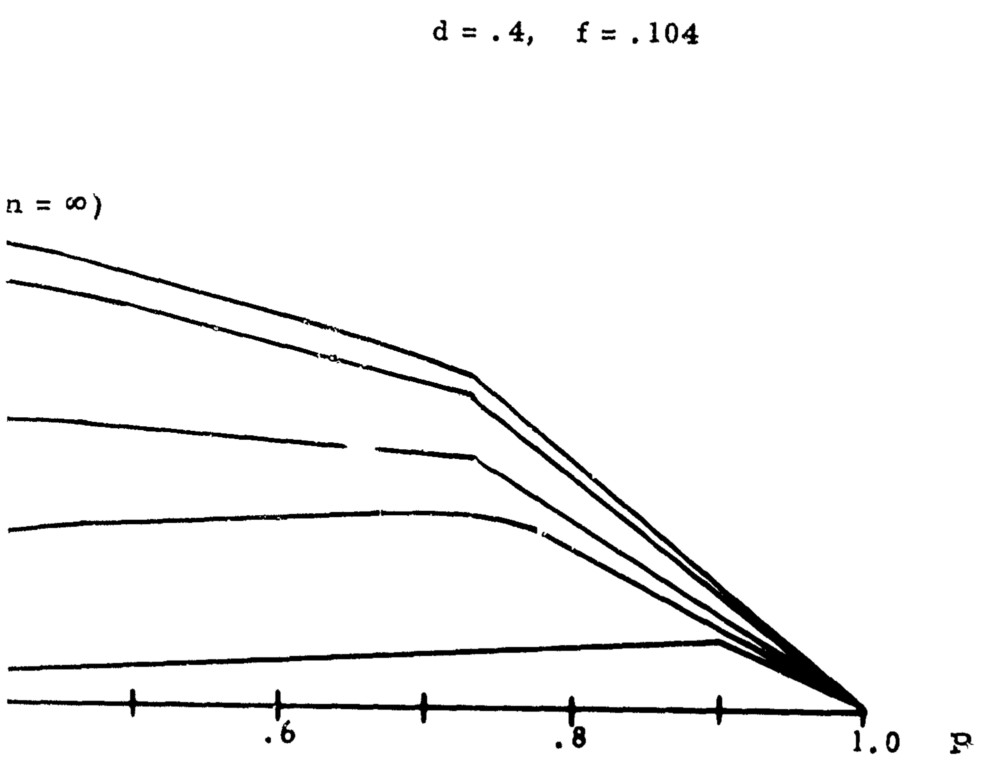

ıvergence of $R_n(P)$ to $R(P)$

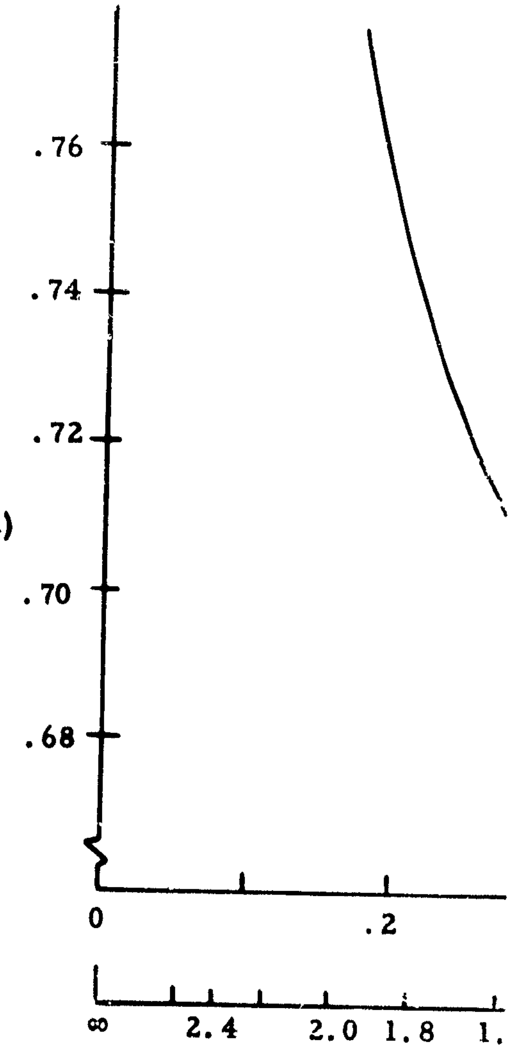

Fi

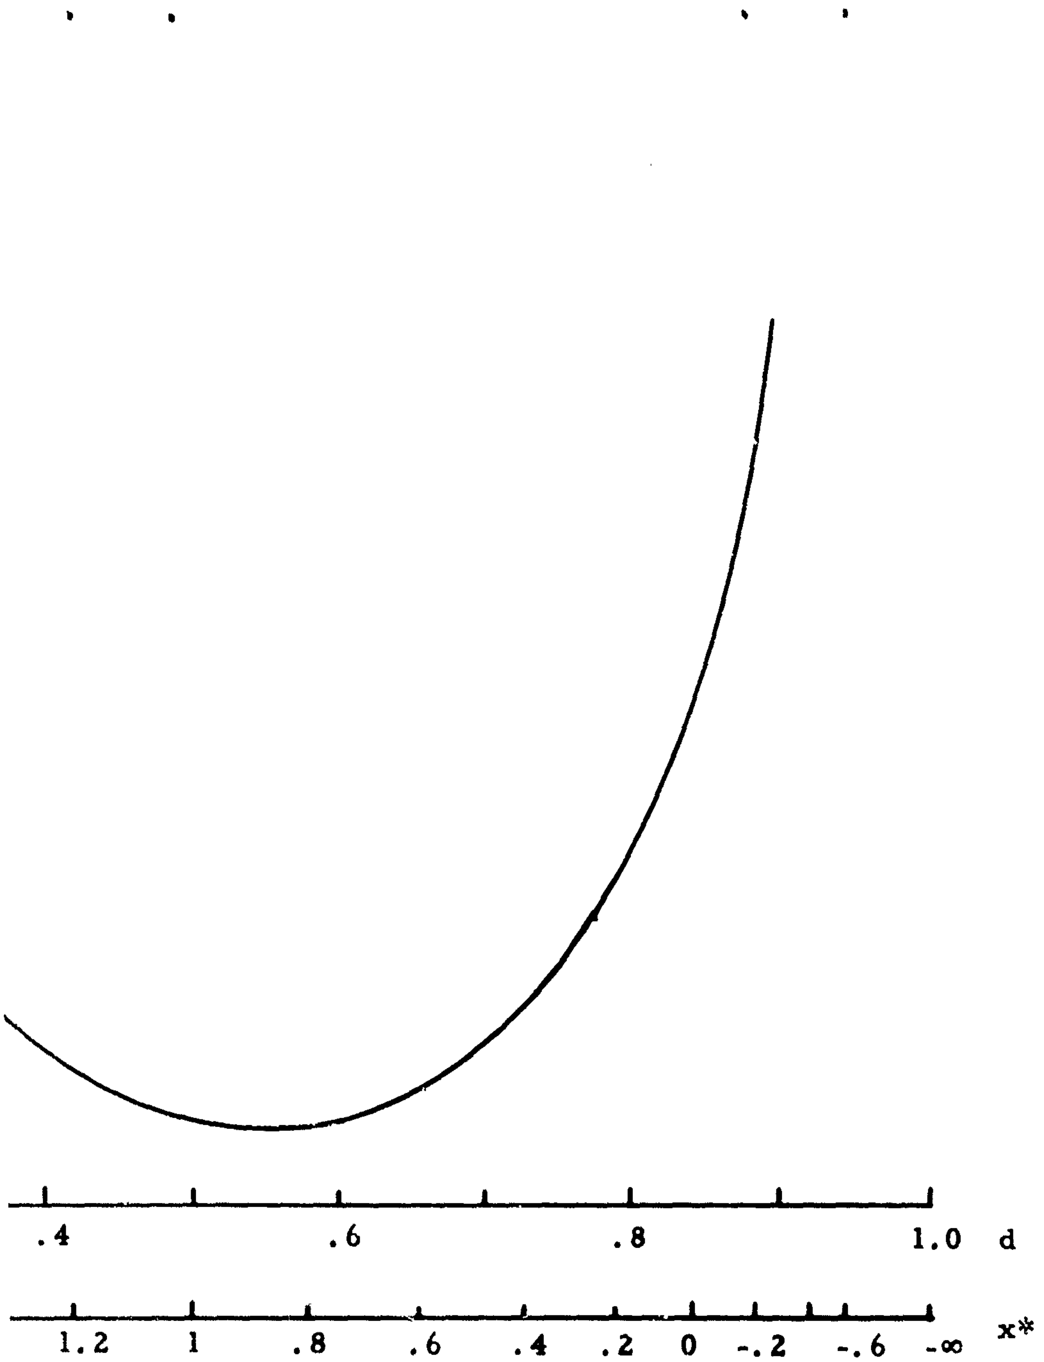

R(λ) as a function of x*

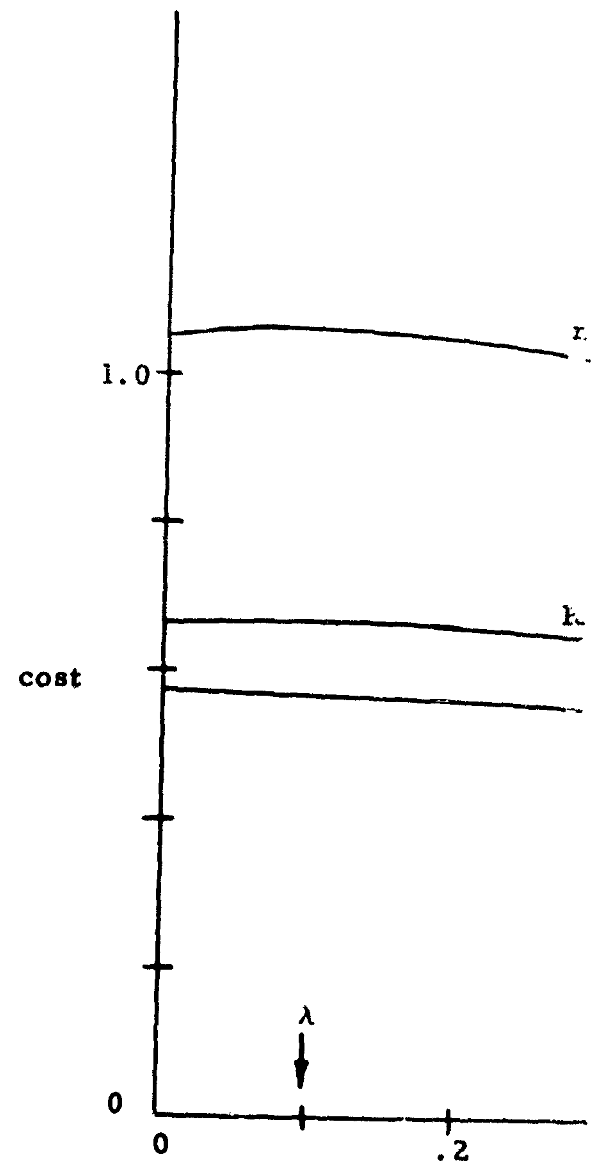

Figure 4.6

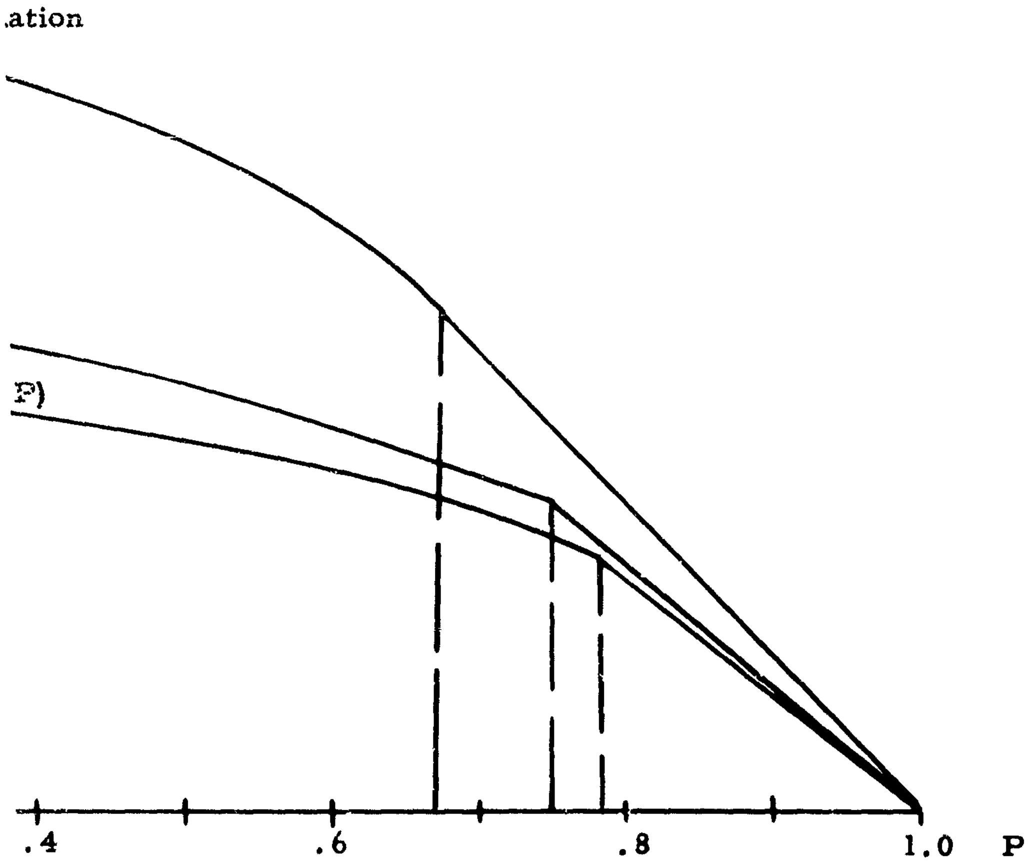

le of $R_{min}(P)$, compared with $V(P)$ with
hout information

Since both these values are less than $\gamma = .75$, another measureme should be taken. Now the a posteriori probability depends upon the sequence y_1, y_2, and is

$$P = \begin{cases} \dfrac{(.15)(1-d) + (.1)(.85)(1-f)}{(.15)(1-d) + (.85)(1-f)} = .18 & (y_1, y_2) = (0,0) \\[2ex] \dfrac{(.15)d + (.1)(.85)f}{(.15)d + (.85)f} = .40 & (y_1, y_2) = (0,1) \\[2ex] \dfrac{(.32)(1-d) + (.1)(.68)(1-f)}{(.32)(1-d) + (.68)(1-f)} = .28 & (y_1, y_2) = (1,0) \\[2ex] \dfrac{(.32)d + (.1)(.68)f}{(.32)d + (.68)f} = 61 & (y_1, y_2) = (1,1) \end{cases}$$

which are all still less than .75. However, we see that the observ with $y_2 = 1$ (i.e. (0,1) and (1,1)) have a higher P. We would ex that in these cases a $y_3 = 1$ observation would result in $P > .75$, in fact this is so. Thus the sequential strategy develops sequences o observations which lead eventually to the decision to send the aircra

4.9 Analysis of Some Non-Sequential and Non-Optimal Decisio Rules: Practical Considerations

Although the threshold search process of the previous sect involves a simpler equation than the one where continuous x is mea the implementation still requires calculation of a weighted sum of li lihood ratios and comparison with a threshold, just as in the earlier sections. For this reason, the required decision making demands a sophisticated arithmetic capability, a condition often lacking in pract situations.

A study of some non-sequential (and therefore non-optimal) decision rules that have the advantage of being simple to use are now considered using the model of a randomly arriving target derived in this chapter. In particular, we would like to evaluate some of the "classical" rules under the condition that, although they were designed for hypothesis testing, they are applied to the randomly arriving target.

The non-sequential STSD rules of section (2.4) are of particular interest. Suppose that they are used in repeated situations, so that at every time interval decision D_1 or D_0 is made, but that the probability P of "target present" is always the same at each situation. We consider the search to continue until the correct decision $\{D_1 | S_1\}$ is made, at which time it terminates. If we wish to consider this within the structure of the randomly arriving target model, we see that it is the equivalent of assuming the rule:

$$\text{for any } x_i \text{, if } x_i \geq x^* : D$$
$$x_i \leq x^* : W$$

or, in words, take action on the first measurement that exceeds some fixed threshold x^*. Let us investigate how x^* is obtained.

With STSD, $C(P)$, the cost <u>per decision</u> is minimized, to get x^*, where

$$C(P) = (1-P)C_{10}f + PC_{01}(1-d) \qquad (4.27)$$

(from equation (2.4), with $f \equiv p_f$, $d \equiv p_d$, $C_{00} = C_{11} = 0$).

The cost structure of the randomly arriving target, however, requires a minimization of the cost of the entire search $V(P)$. The cost of search, given that the target arrives at time t and using the

above strategy, consists of two terms. The first is the cost of false decisions that occur until t, an average of f per unit time. The second term is due to the expected time $\bar{\tau}$ needed to make the correct n-decision after the target arrives, which is

$$\bar{\tau} = \sum_{\tau=0}^{\infty} d(1-d)^{\tau-1} = \frac{1}{d} \quad .$$

The total cost $V(t)$ given arrival time t is thus

$$V(t) = fFt + \frac{W}{d} \quad ,$$

and the average total cost V is

$$V = \sum_{t=0}^{\infty} V(t)\, p(t) = fF\bar{t} + \frac{W}{d} \quad ,$$

where $p(t)$ is the arrival time distribution. By saying that P is the same for each observation we have assumed that $p(t) = P(1-P)^t$, so that $\bar{t} = \frac{1-P}{P}$ and

$$V(P) = V = fF\left(\frac{1-P}{P}\right) + \frac{W}{d} \quad . \tag{4.28}$$

We have seen before that the operating point when the cost is expressed by equation (4.27) is defined by the point on the ROC where

$$\frac{d(d)}{df} = \frac{1-P}{P} \frac{C_{10}}{C_{01}} \quad .$$

By minimizing the cost assumed in equation (4.28), however, we have for an operating point

$$\frac{d(d)}{df} = d^2 \frac{1-P}{P} \frac{F}{W} \quad . \tag{4.29}$$

If we observe that F and W are the exact analogues of the costs C_{10} and C_{01} respectively (that is $C_{10} = F =$ cost of $\{D|S_0\}$, and $C_{01} = W =$ cost of $\{W|S_1\}$) we see that the operating point expressed by equation (4.29) requires a smaller value of $d(d)/df$ on the ROC, and so operation at a higher d and f. See, for example, Figure 2.2.

What is most interesting, however, is that for the randomly arriving target, the two variables of detection that enter into the cost of search are the probability per unit time (or "rate") of making false alarms and the expected time until detection after the target arrives. For the simple decision rule analysed above, the false alarm rate was simply equal to the false indication probability f, since only one indication was required for a D-decision. Similarly, $\bar{\tau}$, the average time until detection after the target arrives is simply related to d since τ, the time until detection after the target arrives is geometrically distributed with parameter d, and hence $\bar{\tau} = (\frac{1}{d})$. In general, however, we see that for the model of the randomly arriving target we can characterise the search system (detection device plus a decision rule) by means of a couple $(\phi, \bar{\tau})$, where ϕ is now defined as the false alarm rate. We recall that a false alarm is the decision to decide the target has arrived when in fact it has not.

This characterisation of a system by $(\phi, \bar{\tau})$ will be shown to be entirely analagous to the characterisation of a detection device by the (p_f, p_d) couple defining the ROC, as in Chapter II.

4.10 The System Operating Characteristic (SOC)

Just as in STSD the receiver operating characteristic (ROC) was a useful concept in comparing and evaluating detection devices, (see section 2.3) a claim is now made for the use of a similar concept

for the comparison and evaluation of certain search systems. By search system we mean a detection device used in conjunction with a decision rule. Although the ROC can be used to compare search systems, it can do so only if they conform to the same set of decision rules, which in turn lead to some equivalent p_f and p_d .

The System Operating Characteristic (SOC) that is proposed here is simply a plot of $\bar{\tau}$ v.s. ϕ for any search system. By ϕ we now mean that fraction of time that the system produces the wrong decision $\{D|S_0\}$, and by $\bar{\tau}$ we mean simply the expected time to reach the decision $\{D|S_1\}$ after the arrival of the target. For the randomly arriving target model the SOC may be used for both qualitative and quantitative comparison of different systems, just as the ROC was used to compare detection devices for the stationary target model.

For example of the qualitative aspects, let us consider two systems S_1 and S_2 with operating points $(\bar{\tau}_1 , \phi_1)$ and $(\bar{\tau}_2 , \phi_2)$. If $\bar{\tau}_1 < \bar{\tau}_2$ and $\phi_1 < \phi_2$, then S_1 is preferable to S_2 , and so the direction of preference on the SOC is "down and to the left". In order to obtain a SOC it is of course necessary to give the decision rule as well as the characteristics of the detection device. The SOC's in Figure 4.7 are the equivalent of the ROC's given in Figure 2.2, i.e. for the detection of a known signal in additive Gaussian noise, (with μ the signal to noise ratio), and the decision rule: if $x \geq x^* : D_1$, if $x \leq x^* : D_0$ at every observation. Note that an increase in μ is still a universal improvement.

Once the SOC is given, the optimal operating point for any decision criterion (not only the "Bayes") may be obtained. As shown in the previous section, the cost of search may be written in general as

$$V = \phi F \bar{t} + W \bar{\tau} \tag{4.30}$$

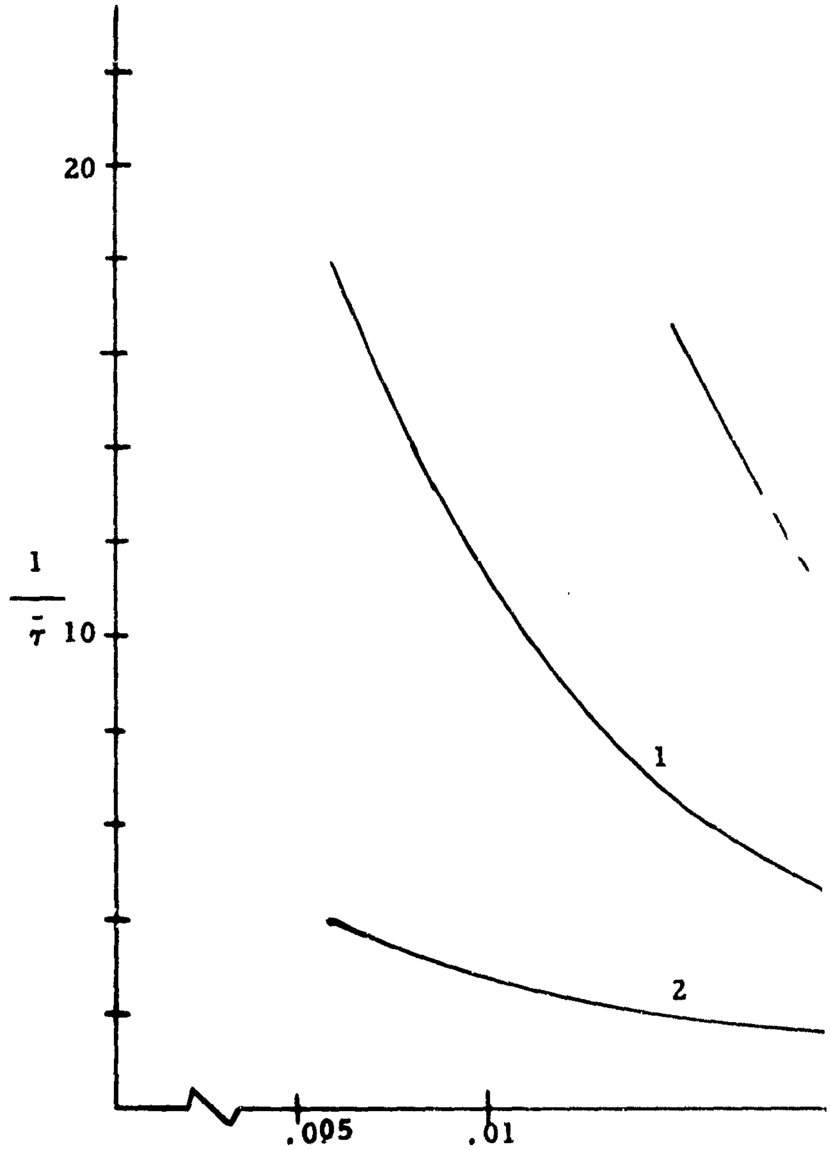

Figure 4.7 System Operati
and device illu

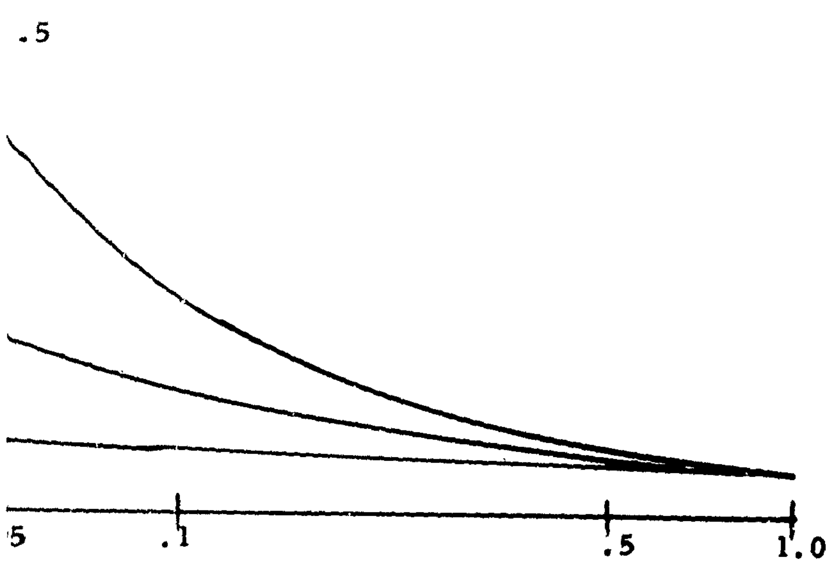

racteristic (SOC) for STSD rule
in Figure 2.1

so that the minimum-cost ("Bayes") operating point is where

$$\frac{d\bar{\tau}}{d\phi} = -\frac{F\bar{t}}{W}$$

with $\bar{t}$ the mean arrival time of the target.

If the Neyman-Pearson criterion is used, some arbitrary maximum allowable false alarm rate $\phi = \phi^*$ is fixed, and the operating point determined by the SOC at the point $\phi = \phi^*$. The use of the SOC for this criterion is very well illustrated by the following example. Suppose that the decision rule is to make decision D only when there are indications $(x \geq x^*)$ at k successive times, and the detection device is for the familiar known target in additive Gaussian noise.

From standard recurrent event theory one can show that

$$\phi = \frac{(1-p_f)\,p_f^{\,k}}{1-p_f^{\,k}}$$

$$\bar{\tau} = \frac{1-p_d^{\,k}}{(1-p_d)\,p_d^{\,k}}$$

where again

$$p_f = \int_{x^*}^{\infty} p_0(x)\,dx$$

$$p_d = \int_{x^*}^{\infty} p_1(x)\,dx$$

For any fixed k and x^* there results a value of ϕ and $\bar{\tau}$, and by letting x^* vary from minus infinity to infinity we produce the SOC for the k-in-a-row system. These SOC's are shown in Figure 4.8 for $k = 1$ (the STSD case), $k = 2$, $k = 3$ with $\mu = 1$. As can be seen, the value of k to be used depends upon where in the $\bar{\tau}$ - ϕ plane the operating point is located.

For the Neyman-Pearson criterion, it is often useful to present ROC as a plot of p_d v.s. μ for the given desired p_f^*. Similarly we may draw the SOC for the above example, for a given ϕ^*, as a plot of $\bar{\tau}$ v.s. μ, or, more convenient, as a plot of $1/\bar{\tau}$ v.s. μ. Figure 4.9 shows such a plot for a fixed $\phi^* = 10^{-8}$. (This value of critical false alarm rate seems to be a popular one for radar search systems.) We note that as μ becomes large $\bar{\tau} \rightarrow k$ (which is reasonable since it still requires k measurements before a D-decision can be made) and the $k = 1$ rule is best. However, as μ becomes small, it becomes advantageous to use the 2-in-a-row rule, then the 3-in-a-row rule, and so on.

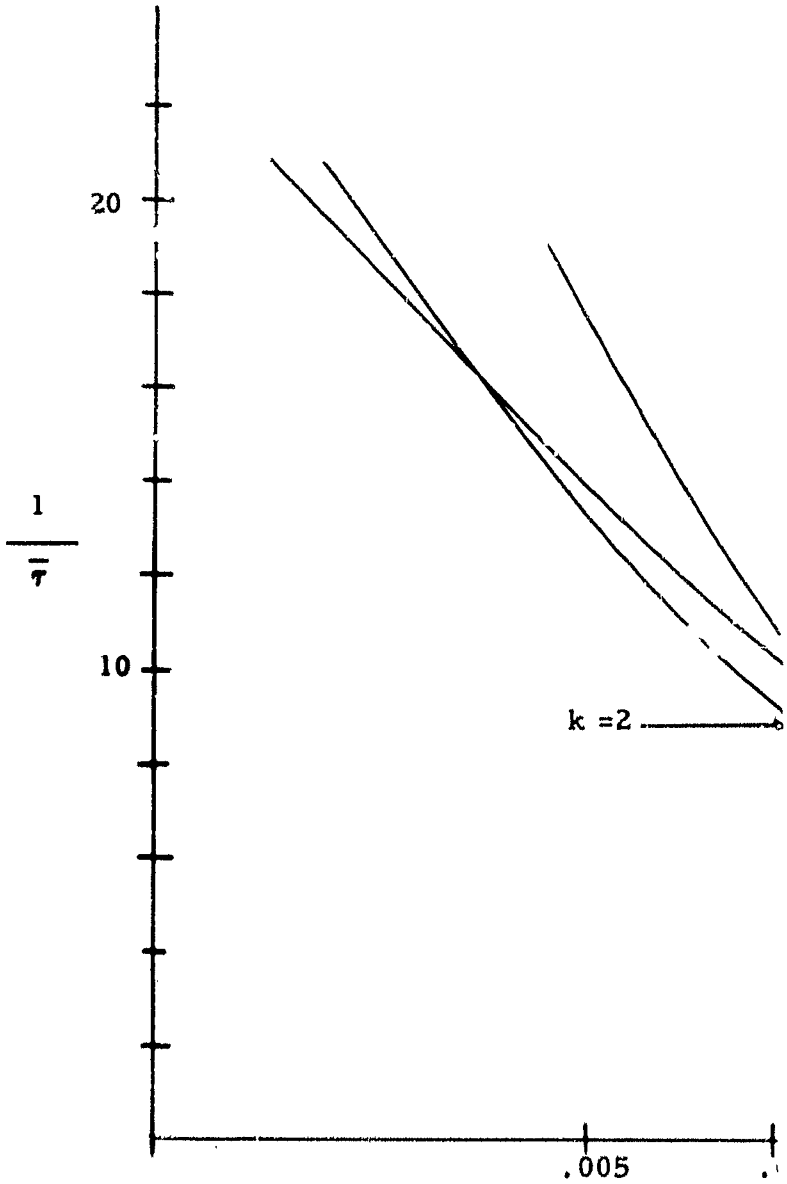

Figure 4.8

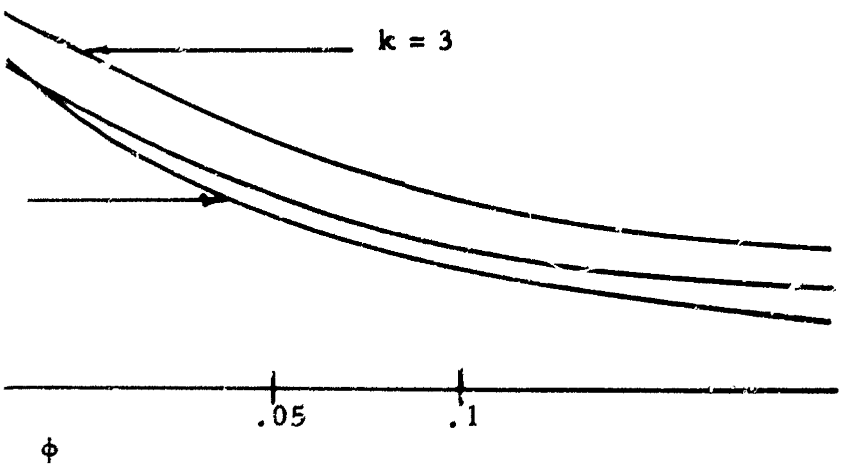

-in-a-row decision rule

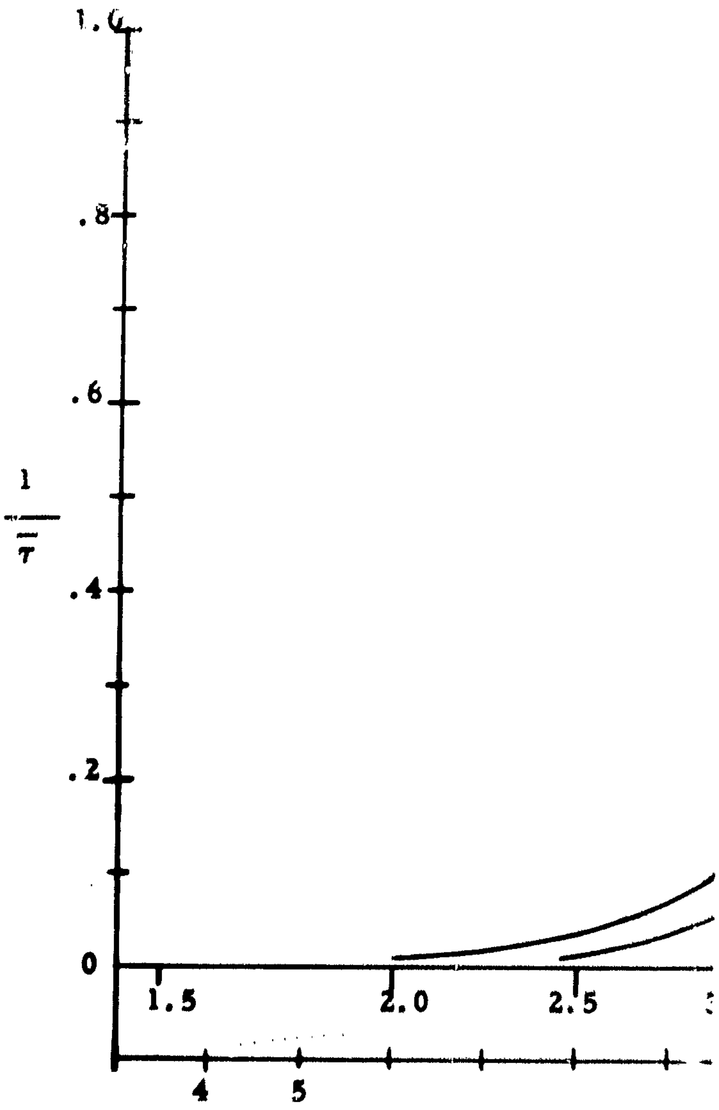

Figure 4

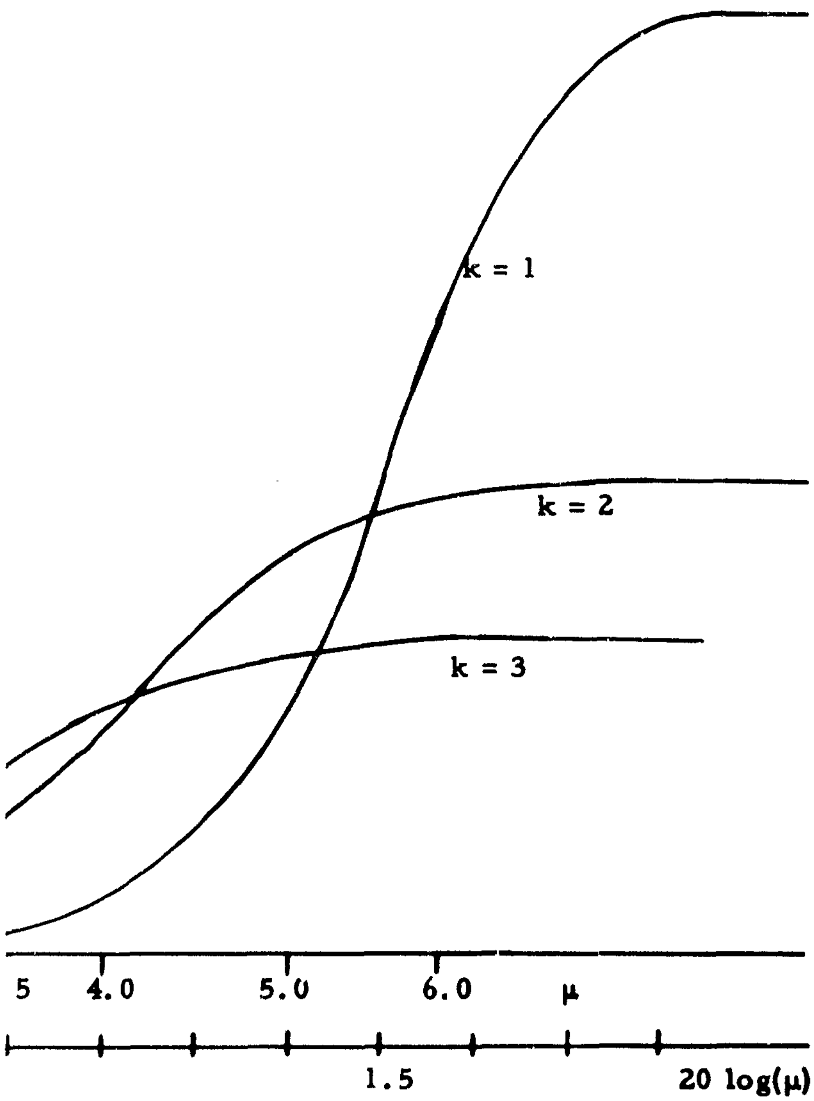

v. s. μ with fixed $\phi^* = 10^{-8}$

CHAPTER V

CONCLUSION

5.1 Summary of Results

The objective of this study has been to obtain a set of decision rules for use in certain general search situations. These rules, called optimal sequential search strategies, are "optimal" in that their use insures that the total cost of search is as small as possible. They are "sequential" in the sense that at any point in the search, the decision rules are affected by what has been observed up to that point.

The principle technique employed has been the representation of the search process in terms of a stochastic dynamic program. This method not only provides the form of the optimal strategy, it also produces as a natural consequence the minimum cost attained by using such a strategy. This allows an immediate quantitative comparison of the results with any non-optimal strategy for which the cost is available.

The models that were treated involved the use of detection devices that are imperfect, in the sense that a target that is present might go undetected, or a "detection" might occur when the target is in fact not present. The latter error is referred to as a false alarm, and its consideration is what distinguishes the present work from most of the published literature in search theory.

5.1.1 Optimal Sequential Search for a Stationary Target

In this model, discussed in Chapter III, the target is considered to be, for all time, present in the region of search with probability P, or not present with probability $(1-P)$. Knowing the value of P, at fixed

intervals of time the searcher must either make a terminal decision or make a measurement of a random variable. The distribution function of this random variable depends upon whether or not the target is actually present. The two terminal decisions involve taking action corresponding to the conclusion that the target is present, or that the target is absent. The search is then ended, and a cost is incurred if the conclusion reached is wrong. If a measurement is taken, a cost (which depends upon whether or not the target is present) is incurred. The measured random variable, however, also depends upon whether or not the target is present, so information about the target's presence is gained from every measurement

The optimal strategy (which decisions to make on the basis of a knowledge of P at the start, plus possible subsequent measurements) and minimum cost were then obtained by the solution of a functional equation. This functional equation may be solved numerically by a method of successive approximations which is equivalent to treating the equation as a dynamic program. The state variable of this program is the probability that the target is present, which is adjusted by Bayes' Rule after every observation.

Proof of the convergence of such a program, and the general form of the strategy that arises were developed and examples calculated for interesting cases. In particular, it was shown that the optimal sequential strategy is similar to the Wald sequential probability ratio test (sprt). The sprt, however, involves the α and β errors of classical statistics and may be only approximately derived, while the new result offers the minimum cost directly, and provides a solution to any desired degree of accuracy. In addition, if the search must be truncated (stopped after some given length of time) the dynamic programming solution still

offers, by definition, the minimum cost search, whereas sprt truncation solutions often are quite arbitrary. The results were also compared with the minimum cost non-sequential decision rule that is a part of the Statistical Theory of Signal Detection.

The method of solution also facilitates the evaluation of certain non-optimal decision rules. In particular, we have analysed the situation where observation of a continuous random variable is restricted to noting whether or not it exceeds a fixed threshold. The extra cost due to this quantization of the observation was then obtained.

The method developed here for treating this problem is in fact applicable to any hypothesis testing situation, where the cost of experimentation depends upon the state of nature, and the a priori hypothesis probability and the terminal decision error costs are known.

5.1.2 Optimal Sequential Search for a Target Arriving at a Random Time

Chapter IV treated a search situation that is often referred to as the "raid recognition" problem. At the start of the search, the target is present in the region of interest with an a priori probability P. At successive constant time intervals thereafter, if the target is not yet present, it has probability λ of arriving. Once the target arrives in the region, it remains for the rest of the search. At every time interval, the searcher must either decide to take action commensurate with the conclusion that the target has arrived, or make an observation of a random variable which has a distribution that depends upon whether or not the target has in fact arrived. If it is concluded that the target has arrived, and it hasn't, the searcher is so informed

and a false alarm cost is incurred. After the target has arrived, a cost is incurred that is proprotional to the length of time it takes to conclude that it has arrived.

The optimal sequential strategy and resulting minimum cost were again obtained from the solution of a functional equation of the dynamic programming type. The state variable of the program is the probability that the target has already arrived. This is adjusted by Bayes' Rule after every observation, or set to zero if a "target has arrived" conclusion is made and the target has in fact not yet arrived.

Comparison between the optimal sequential strategy and some non-optimal strategies was made, and the difference between them discussed.

A side result of the solution has been the development of what is called a System Operating Characteristic. This is simply a plot of time until detection (given the target has arrived) against false alarm rate for any search system (search strategy plus detection device). Its use is analogous in the randomly arriving target search to the use of the Receiver Operating Characteristic in the hypothesis test search, in that it allows both a qualitative and quantitive comparison of various detection devices and decision rules.

It has also been noted that the solution is valid for a larger class of problems than simply those of search. For example, the method developed here may be used to derive optimal checking strategies for machinery subject to random failure while some pertinent output parameter (a random variable) is being monitored.

5.1.3 Comments on the Solutions---Need for Further Analysis

The major results of this study have been the formulation of the search models in a way that takes into account their intrinsic sequential qualities. Due to the dynamic program form of the resulting equations, analytic solutions (except for the non-informative random target arrival time search) are not attainable. An additional analytical difficulty has been the existence of the (reasonable) assumption that the observed random variables are normally distributed. Although the solutions can be obtained numerically for any set of parameters, it would be interesting to obtain at least approximate solutions for certain limiting cases. There are two of these in particular that the author has attempted, with (to date) little results.

The first concerns the limiting behavior of $C(P)$, the minimum cost of sequential search for the stationary target. In many practical cases, P is very close to zero, while the false alarm cost is very much greater than the missed target cost $(C_{10} \gg C_{01})$. How does $C(P)$ vary with μ in this limit, and how does the strategy vary?

The second problem has to do with the very slow convergence of the calculation of $V(P)$, the minimum cost for the randomly arriving target search, when λ, the probability of arrival per unit time, becomes small. If it were possible to start the iteration at some reasonable guess for $V(P)$, the convergence would become that much quicker. The problem of interest, then, is to obtain some reasonable approximation for $V(P)$ as a function of μ as $\lambda \to 0$. In particular, an approximation for $V(0) = V(\lambda)$ would be useful.

In general, any results that can produce an analytic expression for $V(P)$ and $C(P)$, or the decision regions (γ, δ) and $(0, \gamma)$ defining the strategies, for any (non-degenerate) limiting case, would be worthwhile.

5.2 Suggested Areas of Further Research

The subject of sequential estimation as an element of search was quickly mentioned and dropped in Chapter II. This very difficult topic is related to the problem of parameterization of Markov processes, one that has been recently attacked by Kramer (28). Using his results, it might be possible to analyse the following sort of problem. A stationary target is being searched for as in Chapter III. The signal-to-noise ratio μ is uncertain, however, and is in fact a random variable with (say) a known p.d.f. $g(\mu)$. If a decision is made that the target is present, then an estimate of μ must accompany this decision. An appropriate cost of wrong estimation is assumed. What is the best sequential strategy including a rule for estimating μ? This problem could be set up in a dynamic programming fashion if the posteriori p.d.f. on μ, given an observation x, has the same functional form as $g(\mu)$. However, this rarely is the case, and approximate techniques (such as those developed by Kramer) are needed. It is interesting to note that some approximate sequential estimation techniques, based upon fiducial probability arguments, have been derived for this sort of problem (see Turner (44)).

Although the Randomly Arriving Target model was created to represent a more realistic type of search, it too has a basically weak assumption---that after arriving the target remains in the region for the rest of the search. In fact, in many realistic search situations the

target is continuously appearing and dissapearing, often "at random". Kimball (25) and others have formulated models reflecting this effect, but there has been no analysis taking into consideration false alarms. A related problem has been studied by Drake (12), with the emphasis on the information-theoretical aspects of the "noisy" observations of such an appearing-disappearing model. A decision-theory approach to this sort of situation would prove fruitful in such fields as submarine search, detection of epidemics, etc.

A final suggestion for further work concerns what is probably the most well known of results in search theory, Koopman's (26) solution for the distribution of search effort over a continuous field. Here, the target location is represented by a p.d.f. (on a line, say), and the searcher must allocate a fixed quantity of search effort along the line such that the overall probability of detection of the target is a maximum. The detection probability at any point x is assumed to be an increasing function of the search effort placed between x and $x + dx$. Suppose now that there is in addition a false alarm probability at every x that is independent of the target behavior, but a function of the search effort at x. With an appropriate cost structure, what should the distribution of search effort be? The author has pointed out in an earlier paper (37) that the no-false-alarm solution to the discrete-cell search approaches Koopman's solution in the limit. It would be interesting to see the relation between the results in the present work and a solution to the problem above.

APPENDICES

A. Random Walk With Absorbing Barriers

We mention here some results concerning the one-dimensional random walk with absorbing barriers. In section 2.5 we noted the relation between the sprt and this problem. Specifically, if we let z_i be the logarithm of the likelihood ratio $L(x_i)$, and

$$Z_k = \sum_{i=1}^{k} z_i \ ,$$

then the test comparing Z_k to the boundaries a and b is such a random walk. . A solution to this problem should consist of (at least) the probability that the walk ends with absorption at each of the boundaries a and b , and the expected length of time (number of samples) needed to do so, under each state of nature S_0 and S_1 .

Wald (47) has shown the following results that lead to a partial solution. Let the z_i all have the same p.d.f. $g_h(z)$ under the state of nature $S_h (h=0, 1)$. We define $E_h(\xi)$ to be the expected value of random variable ξ under S_h . Then if $\phi_h(t) \equiv E_h(e^{tz})$ is the moment-generating function of $g_h(\cdot)$, and $E_h(z) \neq 0$ and $g_h(z) > 0$ for some $z > 0$ and some $z < 0$, then there exists a $t = t_h$ such that

$$\phi_h(t_h) = 1 \qquad \text{(A. 1)}$$

and the following fundamental identity may be proven

$$E_h\{e^{t Z_n}[\phi_h(t)]^{-n}\} = 1 \qquad \text{(A. 2)}$$

where n (a random variable) is the length of time until absorption.

Now let us define

$$\Pi_h(b, a) = \text{prob.}\,\{Z_n \leq b\}$$

$$1 - \Pi_h(b, a) = \text{prob.}\,\{Z_n \geq a\} \qquad (A.3)$$

where Z_n is the value of Z at termination, and we have taken advantage of the fact that the process terminates with probability one.

We can now use (A.1) and our knowledge of $g_n(z)$ to calculate t_h. By letting $t = t_h$ in (A.2) we also obtain

$$E_h(e^{t_h Z_n}) = 1 \quad . \qquad (A.4)$$

Now let us suppose that when the walk terminates, it does so at exactly $Z_n = a$ or $Z_n = b$ (i.e. we neglect any excess of Z_n over these boundaries). Then we may write (A.4) with the aid of (A.3) as

$$\Pi_h(b, a)\, e^{bt_h} + [1 - \Pi_h(b, a)]\, e^{at_h} = 1$$

or, solving for $\Pi_h(b, a)$

$$\Pi_h(b, a) = \frac{e^{(a-b)t_h} - e^{-bt_h}}{e^{(a-b)t_h} - 1} \qquad (A.5)$$

Using the same approximation we may also write

$$E_h(Z_n) = b\,\Pi_h(b, a) + a\,[1 - \Pi_h(b, a)]$$

and we may easily show that

$$E_h(n) = \frac{E_h(Z_n)}{E_h(z)} \tag{A.6}$$

Identifying $\Pi_0(b,a)$ with $1-p_f$ and $\Pi_1(b,a)$ with $1-p_d$ leads directly to equations (2.13).

The above formulae hold, unfortunately, only for the restriction used above---that the walk ends only with a jump exactly onto the boundary. When $(a-b)$ is very large compared to the values of z_i, then we see that this assumption is reasonable. However in the general case considered in Chapter III we cannot guarantee this because we cannot tell beforehand the values that these boundaries will assume.

Perhaps a more direct way of deriving the $\Pi_h(b,a)$ is to use a Chapman-Kolmogorov equation to describe the walk in going from one step to the next. Thus

$$\Pi_h(b,a) = \int_{-\infty}^{b} g_h(z)\,dz + \int_{b}^{a} g_h(z)\,\Pi_h(b-z, a-z)\,dz \ .$$

The solution for this equation is not, in general known. A most exhaustive and interesting study of such problems has been recently undertaken by J. H. B. Kemperman in The Passage Problem for a Stationary Markov Chain, University of Chicago Press (1961).

B. Decision Regions For C(P) When W_0 or W_1 Are Zero

Theorem D has been proven for the case where the experimental costs W_0 and W_1 are non-zero. We demonstrate here the proof for the condition $W_0 = 0$ or $W_1 = 0$.

1. Let us take $W_0 = 0$, $W_1 > 0$. Then by step 2 of theorem D we still have $\delta_n \leq \Delta < 1$.

2. $$G_n(P) = PW_1 + \int_{-\infty}^{\infty} g(x)\, C_{n-1}\left(\frac{p_1(x)\,P}{g(x)}\right) dx \geq$$
$$\geq PW_1 + \int_{-\infty}^{\infty} g(x)\, C_{n-1}\left(\frac{p_1(x)\,P}{g(x)}\right) dx$$

   for any $y > -\infty$.

3. (2) above becomes a strict inequality if $\int_{-\infty}^{y} g(x)\, dx > 0$. Let us assume this is so.

4. Let y be the solution to
$$\frac{p_1(y)\,P}{p_1(y)\,P + p_0(y)(1-P)} = \delta_{n-1}$$
   so that $y = y(P, \delta_{n-1})$.

5. Let $\gamma_n = P$ in (2) above. By (4) and the fact that $G_n(P) = \gamma_n C_{01}$, we get
$$\gamma_n C_{01} > \gamma_n W_1 + (1-\gamma_n) C_{10} \int_{y(\gamma_n, \delta_{n-1})}^{\infty} p_0(x)\, dx$$

6. Since $\delta_{n-1} \leq \Delta$ by theorem D, (5) becomes
$$\gamma_n > \frac{C_{10} \int_{y(\gamma_n, \Delta)}^{\infty} p_0(x)\, dx}{C_{01} - W_1 + C_{10} \int_{y(\gamma_n, \Delta)}^{\infty} p_0(x)\, dx}$$

7. We postulate that $\gamma_n > \Gamma' \geq 0$. Then

$$\infty \geq y(\Gamma', \Delta) > (\gamma_n, \Delta)$$

and (6) becomes

$$\gamma_n > \frac{C_{10} \int_{y(\Gamma', \Delta)}^{\infty} p_0(x)\, dx}{C_{01} - W_1 + C_{10} \int_{y(\Gamma', \Delta)}^{\infty} p_0(x)\, dx}$$

8. By letting the right-hand side of (7) be equal to Γ' , we note that $\Gamma' = 0$ is a solution, and we have proved that $\lim_{n \to \infty} \gamma_n = \gamma > 0$.

If the assumption made in (3) does not hold then the proof is invalid, and in fact $\gamma = 0$. A similar proof is followed for the condition $W_0 > 0$, $W_1 = 0$.

C. Flow Chart of Computer Program

We present in Figure C.1 the flow chart of the computer program that iteratively calculates $C(P)$ by use of equation (3.3). The programs used to evaluate other quantities in this work are similar in structure and so are not shown.

The x-axis (where x is the observed random variable) extends from -4 to +12 and is represented by 161 points .1 apart. The P-axis goes from 0 to 1 and is represented by 101 points .01 apart. It is assumed that $f_N(x; 0, 1) = 0$ for $x < -4$ and $x > 4$.

C_0 is the initial value for the iteration and is read in for each calculation. To evaluate $C(P)$ when $(.01)n < P < (.01)(n+1)$, where $n = 0, 1, \ldots, 99$, the linear approximation

$$C(P) \cong C[(.01)n]\,(n+1-100P) \doteq C[(.01)(n+1)]\,(n-100P)$$

is used, which is exact when $C(P) = T(P)$.

The integral

$$\int_{-\infty}^{\infty} g(x)\, C_{n-1}\left(\frac{p_1(x)\,P}{g(x)}\right) dx$$

is approximated by the sum

$$(.1) \sum_{x_i=-4}^{12} g(x_i)\, C_{n-1}\left(\frac{p_1(x_i)\,P}{g(x_i)}\right)$$

which produces for

$$\int_{-\infty}^{\infty} p_0(x)\, dx = \int_{-\infty}^{\infty} p_1(x)\, dx$$

the value .9994.

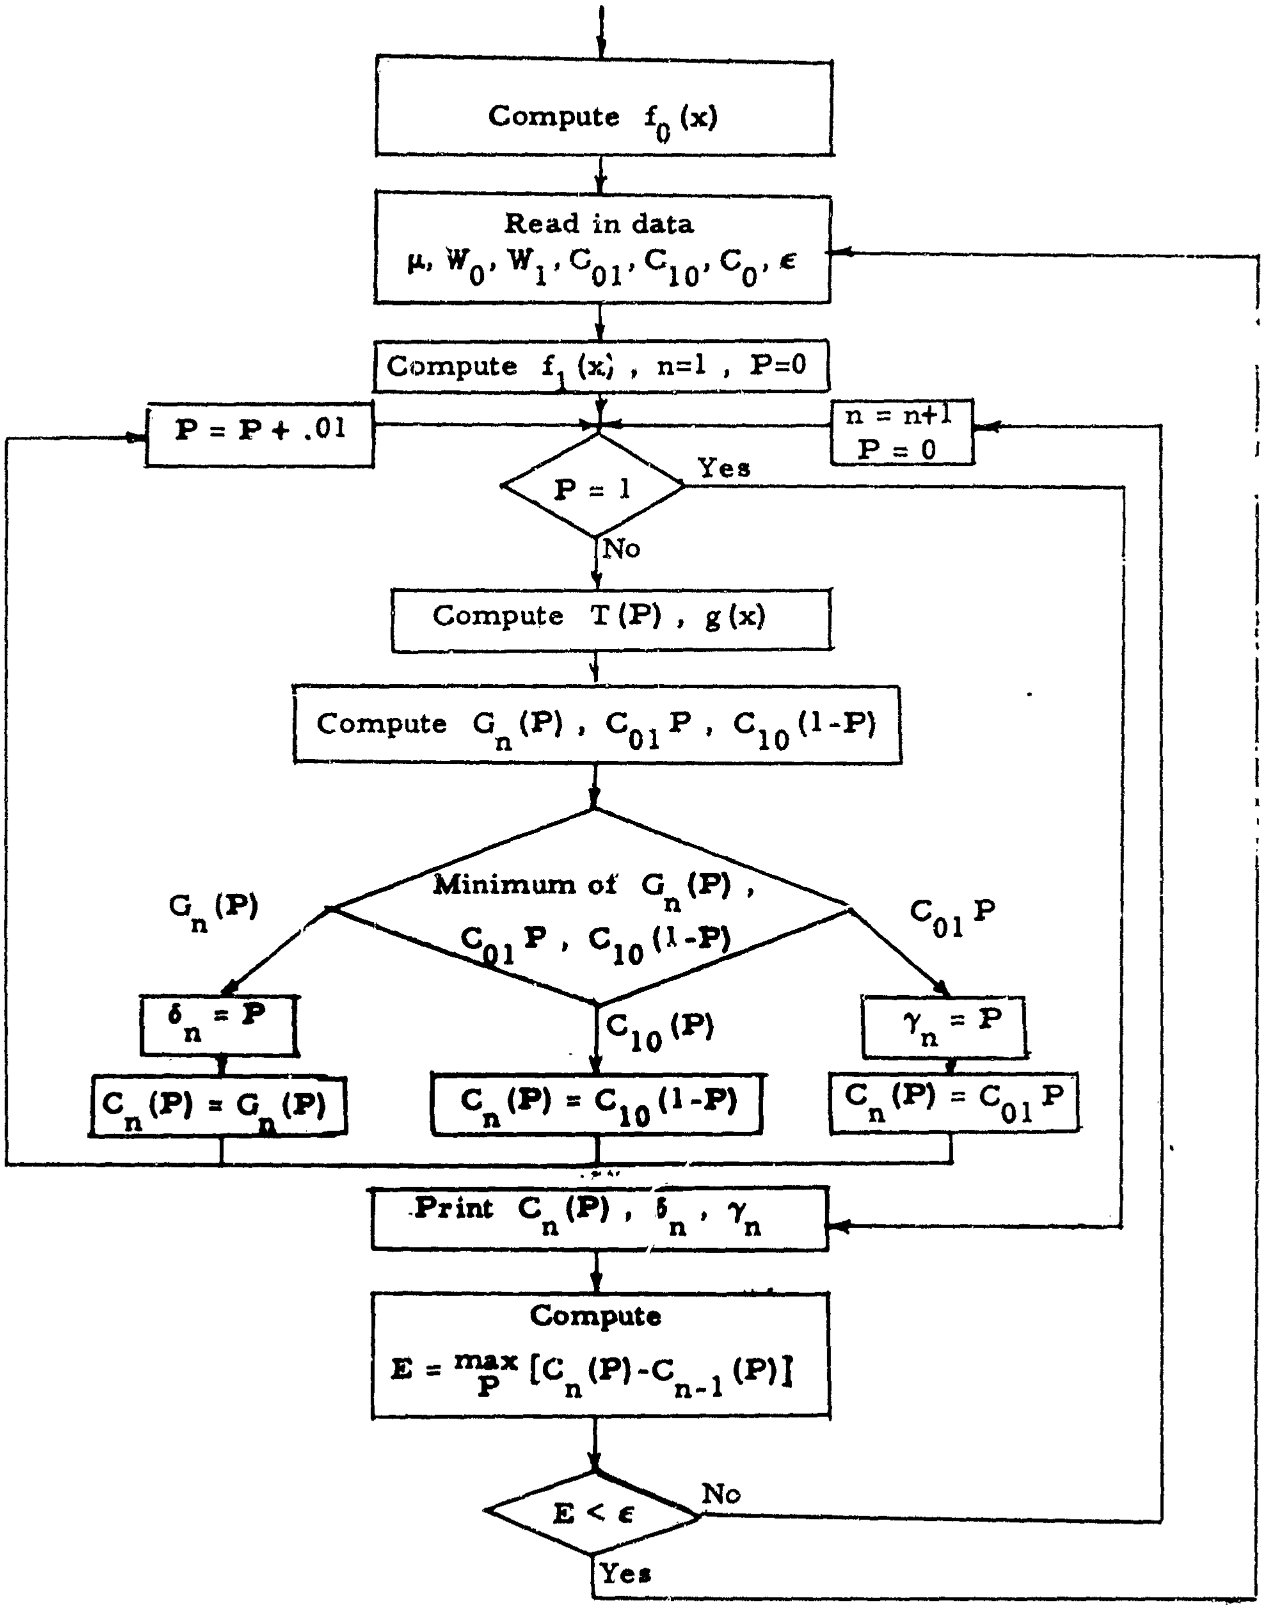

FIGURE C.1

Flow Chart for Computation of C(P)

BIBLIOGRAPHY

1. Banta, E. D.; "On the Detection of Maneuvering Targets"; General Atronics Corp., Rep. #943-223-41 (1961).

2. Barlow, R. E., Hunter, L. C., and Proschan, F.; "Optimum Checking Procedures", S.I.A.M. 11, 1078-1095 (1963).

3. Bellman, R.; Dynamic Programming; Princeton University Press (1957).

4. Bellman, R., Kalaba, R., and Middleton, D.; "Dynamic Programming, Sequential Estimation, and Sequential Detection Processes"; Proc. N.A.S., 47, 338-341 (1961).

5. Blackman, N. M., and Proschan, F.; "Optimum Search for Objects Having Unknown Arrival Times"; J.O.R.S.A. 7, 625-638 (1959).

6. Blackman, N. M.; "Prolegomena to Optimal Discrete Search Procedures", N.R.L.Q. 6, 273-281 (1959).

7. Blackwell, D., and Girschick, M. A.; Theory of Games and Statistical Decisions; Wiley, New York (1954).

8. Bussgang, J. J.; Sequential Detection of Signals in Noise, Doctoral Dissertation; Harvard University (1955).

9. Bussgang, J. J., and Middleton, D.; "Optimum Sequential Detection of Signals in Noise"; IRE Transactions on Information Theory; IT-1 (1955); IT-3 (1954).

10. Charnes, A., and Cooper, W. W.; "The Theory of Search: Optimum Distribution of Research Effort"; M.S. 5, 44-50 (1958).

11. Danskin, J. M.; "Solution of a Search Problem Against Space Probes"; (Unpublished, undated).

12. Drake, A. W.; "Observation of a Markov Process Through a Noisy Channel"; D.Sc. Thesis, Dept. of E.E., M.I.T. (1962).

13. Gilbert, E. N.; "Optimal Search Strategies"; S.I.A.M. 7, 413-424 (1959).

14. Gluss, B.; "An Alternative Solution to the 'Lost at Sea' Problem"; N.R.L.Q. 8, 117-122 (1961).

15. Gluss, B.; "The Minimax Path in a Search for a Circle in a Plane"; N.R.L.Q. 8, 357-360 (1961).

16. Goode, H. H.; "Deferred Decision Theory", Recent Advances in Information and Decision Processes, Ed. Machol, R.; and Gray, P.; Macmillan, N.Y. (1962).

17. Gray, K. B.; "A Condition for Optimality in Sequential Signal Detection", Western Management Science Institute, Working Paper #46 (1963).

18. DeGuenin, J.; "Optimum Distribution of Effort: An Extension of the Koopman Basic Theory"; J.O.R.S.A., 9, 1-7 (1961).

19. Helstrom, C. W.; Statistical Theory of Signal Detection; Pergamon Press (1960).

20. Helstrom, C. W.; "A Range-Sampled Sequential Detection System", P.G.I.T. 3, 43-47 (1962).

21. Hofstetter, E. M.; "Large Sample Sequential Decision Theory"; T. R. #359, R.L.E., M.I.T. (1957).

22. Isbell, J. R.; "An Optimal Search Pattern"; N.R.L.Q. 4, 357-360 (1957).

23. Jeffrey, H.; Theory of Probability (2nd edition), Clarendon Press, Oxford (1948).

24. Kennedy, R. S.; "A Theory of Multiple-Scan Surveillance Systems"; Servomechanism Lab.; M.I.T. (1959).

25. Kimball; "Detection and Tracking as a Markov Process"; NATO A.S.W. Meeting, Siena, Italy (1963).

26. Koopman, B. O.; "The Theory of Search: Part I, Kinematic Bases"; J.O.R.S.A. 4, 324-346 (1956); "The Theory of Search: Part II, Target Detection"; J.O.R.S.A. 4, 503-532 (1956); "The Theory of Search, Part III: The Optimum Distribution of Searching Effort"; J.O.R.S.A. 5, 613-626 (1957).

27. Koopman, B. O.; O.E.G. Report No. 56, "Search and Screening" (1946).

28. Kramer, J. D. R. Jr.; "Partially Observable Markov Processes"; D.Sc. Thesis, Dept. of E.E., M.I.T. (1964).

29. Luce, D., and Raiffa, A. H.; Games and Decisions, Wiley, New York (1957).

30. MacQueen, J., and Miller, R. J.; "Optimal Persistence Policies"; J.O.R.S.A. 8, 362-380 (1960).

31. Marcum, J. I., and Swerling, P.; "Sequential Detection in Radars"; P.G.I.T. 8, 237-245 (1962).

32. Mela, D. F.; "Information Theory and Search Theory as Special Classes of Decision Theory"; J.O.R.S.A. 9, 907-909 (1961).

33. Middleton, D.; An Introduction to Statistical Communication Theory; McGraw Hill, New York (1960).

34. Middleton, D., and Van Meter, D.; "Detection and Extraction of Signals From Noise From the Point of View of Statistical Decision Theory"; S.I.A.M. 3, 192-253 (1955); S.I.A.M. 4, 86-119 (1956).

35. Neuts, M.; "A Multistage Search Game"; Recent Advances in Game Theory, Princeton University Conference (1962).

36. Norris, R. C.; "Studies in Search for a Conscious Evader"; M.I.T. Lincoln Labs, T. R. #279 (1962).

37. Pollock, S. M.; "Optimal Sequential Strategies for Two Region Search When Effort is Quantized"; I.T.R. #14, M.I.T., O. R. Center (1960).

38. Preston, G. W.; "The Search and Detection Efficiency of Surveillance and Communication Devices Using Sequential Probability Ratio Analysis"; Proc. of Symposium on Decision Theory and Applications to Electronic Equipment Development (1960).

39. Reed, J. S.; "Analysis of Signal Detection by Sequential Observer"; M.I.T. Lincoln Labs, T. R. #20 (1953).

40. Schwartz, M.; "A Statistical Approach to the Automatic Search Problem"; Doctoral Dissertation, Harvard University (1951).

41. Selin, I.; "The Sequential Estimation and Detection of Signals in Random Noise"; Rand, RM-2994-PR (1962).

42. Selin, I.; "The Design of a Sequential Test for Detection of Known Signal in Normal Non-White Noise"; Rand, P-2487 (1961).

43. Sponsler, G. C.; "Markov Process Representation of Radar Detection Trials of a Binary Data Sequence"; M.I.T. Lincoln Labs, T. R. # 114.

44. Turner, R. D.; "Operations Research on Recognition. The Generalized Search Process"; Advanced Electronics Center, G.E. Ithaca; Sci. Rep. #2 (1961).

45. Wainstein, L. A.; Zubakov, V. D.; Extraction of Signals from Noise; Prentice Hall (1962).

46. Wald, A.; Statistical Decision Functions; Wiley, New York (1950).

47. Wald, A.; Sequential Analysis; Wiley, New York (1947).

48. Wald, A., and Wolfowitz, J.; "Optimum Character of the Sequential Probability Ratio Test"; Ann. Math. Stat. 19, 326 (1948).

BIOGRAPHICAL SKETCH

Stephen M. Pollock was born in New York City on February 15, 1936. He was raised in Brooklyn, where he attended grade school and James Madison High School. In the fall of 1953 he entered Cornell University and was graduated in 1958 with the degree of Bachelor of Engineering Physics. While an undergraduate he was the recipient of several scholarships, including the New York State Regents Scholarship.

Mr. Pollock entered the Massachusetts Institute of Technology in 1958 and received the degree of Master of Science in 1960. While at M. I. T. he was a Research Assistant in the Physics Department, working at the Operations Research Center and the Synchrotron, except for the academic year 1960-1961 when he was a Space Technology Laboratories Fellow.

The industrial experience of Mr. Pollock has included summer employment at AVCO Research and Advanced Development Division, Space Technology Laboratories, and Arthur D. Little, Inc.

Mr. Pollock is a member of Sigma Xi and the Operations Research Society of America. He is the author of "Optimal Sequential Strategies for Two Region Search When Effort is Quantized", a report of the M. I. T. Operations Research Center. He is also co-author of the paper "Elastic π-p Scattering at 868, 900, 1200 and 1310 Mev." given at the April 1963 Meeting of the American Physical Society.

In 1962 Mr. Pollock married the former Bettina Dorn of Brockton, Massachusetts.